Greedy Sampling of Graph Signals

Luiz F. O. Chamon and Alejandro Ribeiro

arXiv:1704.01223v1 [cs.IT] 5 Apr 2017

***Abstract*—Sampling is a fundamental topic in graph signal processing, having found applications in estimation, clustering, and video compression. In contrast to traditional signal processing, the irregularity of the signal domain makes selecting a sampling set non-trivial and hard to analyze. Indeed, though conditions for graph signal interpolation from noiseless samples exist, they do not lead to a unique sampling set. Thus, the presence of noise makes sampling set selection a hard combinatorial problem. Although greedy sampling schemes have become ubiquitous in practice, they have no performance guarantee. This work takes a twofold approach to address this issue. First, universal performance bounds are derived for the interpolation of stochastic graph signals from noisy samples. In contrast to currently available bounds, they are not restricted to specific sampling schemes and hold for any sampling sets. Second, this paper provides near-optimal guarantees for greedy sampling by introducing the concept of approximate submodularity and updating the classical greedy bound. It then provides explicit bounds on the approximate supermodularity of the interpolation mean-square error showing that it can be optimized with worst-case guarantees using greedy search even though it is not supermodular. Simulations illustrate the derived bound for different graph models and show an application of graph signal sampling to reduce the complexity of kernel principal component analysis.**

***Index Terms*—Graph signal processing, sampling, approximate submodularity, greedy algorithms, kernel multivariate analysis.**

I. Introduction

Graph signal processing (GSP) is an emerging field that studies signals supported on irregular domains [1], [2]. It extends traditional signal processing techniques to more intricate data structures, finding applications in sensor networks, image processing, and clustering, to name a few [3]–[5]. Extensions of sampling, in particular, have attracted considerable interest from the GSP community [6]–[12]. This is not surprising given the fundamental role of sampling in signal processing. Sampling methods in GSP are broadly divided into two categories: *selection sampling*, in which the graph signal is observed at a subset of nodes [11], and *aggregation sampling*, in which the signal is observed at a single node for many applications of the graph shift [7]. This work focuses on the former.

As in classical signal processing, samples are only useful inasmuch as they represent the original signal. Conditions under which it is possible to reconstruct a graph signal from noiseless samples can be found in [8]–[11]. However, because they do not necessarily lead to a unique sampling set, the presence of noise raises the issue of which sampling set to choose. In general, however, selecting an optimal sampling set is NP-hard [13]–[16].

Electrical and Systems Engineering, University of Pennsylvania. e-mail: luizf@seas.upenn.edu, aribeiro@seas.upenn.edu. Part of the results in this paper appeared in [21] and [31].

In [6], [11], this issue was addressed using randomized sampling schemes, for which optimal sampling distributions and performance bounds were derived for different types of graphs and graph signals. It was shown that, for high signal-to-noise ratio (SNR), sampling proportionally to the leverage score (or its square-root) approximates the sampling distribution that minimizes the reconstruction mean-square error (MSE). Alternatively, a convex relaxation approach was adopted in [17], where the sampling set selection problem was cast as a binary semi-definite program (SDP) and solved by relaxing the binary constraint and truncating the solution. Nevertheless, greedy sampling remains ubiquitous and has proven successful in many applications [8]–[11], [17], [18], though performance analyses are available only for surrogate figures of merit of the MSE [12].

To be sure, this is not surprising given the attractive features of greedy algorithms for large-scale problems. First, their complexity is polynomial in the deterministic case and stochastic versions exist that are linear in the size of the ground set [19]. Also, since they build the solution sequentially, they can be interrupted at any time if, for instance, a desired performance level is reached. More importantly, there is an upper bound on the suboptimality of the greedy solution to monotonic supermodular function minimization problems. This is indeed why greedy algorithms are often used in sensor selection, experimental design, and machine learning [13]–[16], [20]. However, the main performance measure in GSP, namely the MSE, is not supermodular in general [21].

In this work, we study the reconstruction (interpolation) performance of greedy sampling schemes in GSP and set out to reconcile the empirical success of greedy MSE minimization and the fact that it is not supermodular. First, in contrast to [6], [11], we model graph signals as random processes and take sampling to be deterministic (Section II). Then, we derive bounds on the interpolation MSE that are universal in the sense that they hold for all sampling sets and any sampling method (Section III). These universal bounds are explicit, tractable, and provide practical means of benchmarking the MSE performance of any sampling scheme. Numerical analyses show that the bounds are tight when signal and noise are homeoscedastic. Finally, we develop the concept of *approximate supermodularity* introduced in [21] and provide near-optimal guarantees for the greedy minimization of the interpolation MSE (Section IV). This result justifies the use of greedy sampling set selection in GSP and explains its success.

To illustrate the practical value of these results, we recall that the concept of sampling is also at the core of statistical methods, such as data subsetting and variable selection, that are crucial for *big data* applications [22], [23]. Kernel methods, in particular, are prone to complexity issues in large data sets. For instance, performing kernel principal component

analysis (kPCA) on a data set of size n requires n^2 kernel evaluations (KEs) and $\Theta(n^3)$ operations and extracting projections for new data takes n KEs and $\Theta(np)$ operations, where p is the number of principal components (PCs) retained [24], [25]. We show that this problem can be cast in the context of GSP and that greedy sampling can be used to reduce its complexity at a small performance cost (Section V-B).

Notation: Lowercase boldface letters represent vectors ($\boldsymbol{x}$), uppercase boldface letters are matrices ($\boldsymbol{X}$), and calligraphic letters denote sets ($\mathcal{A}$). We write $|\mathcal{A}|$ for the cardinality of $\mathcal{A}$. Set subscripts refer either to the vector obtained by keeping only the elements with indices in the set ($\boldsymbol{x}_{\mathcal{A}}$) or to the submatrix whose columns have indices in the set ($\boldsymbol{X}_{\mathcal{A}}$). To say $\boldsymbol{X}$ is a positive semi-definite (PSD) matrix we write $\boldsymbol{X} \succeq 0$, so that for $\boldsymbol{X}, \boldsymbol{Y} \in \mathbb{C}^{n\times n}$, $\boldsymbol{X} \preceq \boldsymbol{Y} \Leftrightarrow \boldsymbol{b}^H\boldsymbol{X}\boldsymbol{b} \leq \boldsymbol{b}^H\boldsymbol{Y}\boldsymbol{b}$, for all $\boldsymbol{b} \in \mathbb{C}^n$. The set of PSD matrices is denoted $\mathbb{S}_+$. Finally, we take the derivative of a function f with respect to an $n \times 1$ vector $\boldsymbol{x}$ to yield the $1 \times n$ gradient vector, i.e., $\frac{\partial f}{\partial \boldsymbol{x}} = [\ \partial f/\partial x_1 \ \cdots \ \partial f/\partial x_n\ ]$ [26].

II. Sampling and Interpolation of Graph Signals

A graph-supported signal, or *graph signal* for short, is an assignment of values to the nodes of a graph. Formally, let $\mathbb{G}$ be a weighted graph with node set $\mathcal{V}$, having cardinality $|\mathcal{V}| = n$, and define a graph signal to be an injective mapping $\sigma : \mathcal{V} \to \mathbb{C}$. For an ordering of the nodes in $\mathcal{V}$, this signal can be represented as an $n \times 1$ vector that captures its values at each node:

$$\boldsymbol{x} = [\ \sigma(u_1) \quad \cdots \quad \sigma(u_n)\ ]^T, \quad u_i \in \mathcal{V}. \tag{1}$$

In what follows, we assume that the node ordering is fixed, so that we can index $\boldsymbol{x}$ using elements of $\mathcal{V}$. For instance, we write $\boldsymbol{x}_{\{u_i,u_j,u_k\}} = [\ \sigma(u_i)\ \sigma(u_j)\ \sigma(u_k)\ ]^T$.

Of interest to GSP is the spectral representation of the signal σ (or $\boldsymbol{x}$), which depends on the graph on which it is supported. Indeed, let $\boldsymbol{A} \in \mathbb{C}^{n\times n}$ be a matrix representation of $\mathbb{G}$. Usual choices include the adjacency matrix or one of the discrete Laplacians [1], [2]. Assume that $\boldsymbol{A}$ is consistent with the signal vector (1) in the sense that they employ the same ordering of the nodes in $\mathcal{V}$. Furthermore, assume that $\boldsymbol{A}$ is normal, i.e., that there exist $\boldsymbol{V} \in \mathbb{C}^{n\times n}$ unitary and $\boldsymbol{\Sigma} \in \mathbb{R}^{n\times n}$ diagonal such that $\boldsymbol{A} = \boldsymbol{V}\boldsymbol{\Sigma}\boldsymbol{V}^H$, where $\cdot^H$ denotes the Hermitian (conjugate transpose) operator [27]. Then, the *graph Fourier transform* of $\boldsymbol{x}$ is given by [1], [2]

$$\bar{\boldsymbol{x}} = \boldsymbol{V}^H\boldsymbol{x}. \tag{2}$$

Observe that if $\boldsymbol{A}$ is normal we obtain a spectral energy conservation properties analog to Parseval's theorem in classical signal processing. It is ready to see that $\|\bar{\boldsymbol{x}}\|_2 = \|\boldsymbol{x}\|_2$ if and only if $\boldsymbol{V}$ in (2) is unitary, which holds if and only if $\boldsymbol{A}$ is normal [27].

Similar to traditional signal processing, a graph signal $\boldsymbol{x}$ is said to be bandlimited when its spectral representation is sparse. Explicitly, $\boldsymbol{x}$ is $\mathcal{K}$-*bandlimited* if $\bar{\boldsymbol{x}}$ in (2) is $\mathcal{K}$-sparse, i.e., $\bar{\boldsymbol{x}}_{\mathcal{V}\setminus\mathcal{K}}$ is a zero vector. Then,

$$\boldsymbol{x} = \boldsymbol{V}_{\mathcal{K}}\bar{\boldsymbol{x}}_{\mathcal{K}}. \tag{3}$$

Note that in classical signal processing, a signal is considered bandlimited when its spectral representation has finite support ("low-pass"). Although there have been works exploiting graph frequency orderings based on the eigenvalues of $\boldsymbol{A}$ [6], [10], [11], we abuse the term bandlimited since we do not rely on any such orderings.

The interest in bandlimited graphs signals is motivated similarly to traditional signal processing: these signals can be sampled and interpolated without loss of information. Indeed, take sampling to be the operation of observing the value of a graph signal on $\mathcal{S} \subseteq \mathcal{V}$, the *sampling set*. Then, there exists a set $\mathcal{S}$ of size $|\mathcal{K}|$ such that $\boldsymbol{x}$ can be recovered exactly from $\boldsymbol{x}_{\mathcal{S}}$ [8]–[11]. If, however, only a corrupted version of $\boldsymbol{x}_{\mathcal{S}}$ is available, then $\boldsymbol{x}$ can only be approximated. To do so, the next section poses noisy interpolation as a stochastic estimation problem, from which the minimum MSE interpolation operator can be derived. This then allows us to provide universal bounds on the reconstruction error and give near-optimal guarantees for greedy sampling strategies.

A. Graph signal interpolation

To study the interpolation of graph signals from noisy samples, start by considering the class of bandlimited graph signals that are *wide-sense stationary* random processes with respect to $\mathbb{G}$ [28]–[30]. In other words, take $\bar{\boldsymbol{x}}_{\mathcal{K}}$ in (3) to be a zero-mean circular random vector with $\boldsymbol{\Lambda} = \mathbb{E}\,\bar{\boldsymbol{x}}_{\mathcal{K}}\bar{\boldsymbol{x}}_{\mathcal{K}}^H = \operatorname{diag}\{\lambda_i\}$, $\lambda_i \in \mathbb{R}_+$. Without loss of generality, assume $\boldsymbol{\Lambda}$ is full-rank. Otherwise, remove from $\mathcal{K}$ any element i for which $\lambda_i = 0$. Then, let $\boldsymbol{y}$ be a noisy version of the graph signal of interest, i.e.,

$$\boldsymbol{y} = \boldsymbol{x} + \boldsymbol{w}, \tag{4}$$

where $\boldsymbol{w} \in \mathbb{C}^n$ is a circular zero-mean noise vector with diagonal covariance matrix $\boldsymbol{\Lambda}_w = \mathbb{E}\,\boldsymbol{w}\boldsymbol{w}^H = \operatorname{diag}\{\lambda_{w,i}\}$, $\lambda_{w,i} \in \mathbb{R}_+$ and $\lambda_{w,i} > 0$. Note that (4) is related to the class of *approximately bandlimited* graph signals from [6] and that the noiseless case can be recovered by taking $\boldsymbol{w} = \boldsymbol{0}$.

The signal in (4) is then sampled by observing the elements of $\boldsymbol{y}$ whose index are in the sampling set $\mathcal{S}$. To clarify the derivations, define the selection matrix $\boldsymbol{C} \in \{0,1\}^{|\mathcal{S}|\times N}$ composed of the identity matrix rows with indices in $\mathcal{S}$, so that the samples of (4) can be written as

$$\boldsymbol{y}_{\mathcal{S}} = \boldsymbol{C}\boldsymbol{y}. \tag{5}$$

For generality, suppose that the samples in (5) are used to reconstruct

$$\boldsymbol{z} = \boldsymbol{H}\boldsymbol{x}, \tag{6}$$

for some matrix $\boldsymbol{H} \in \mathbb{C}^{m\times n}$, as opposed to the original graph signal $\boldsymbol{x}$. This accounts for scenarios in which the quantity of interest is a linear transformation of $\boldsymbol{x}$, such as the output of a classifier or estimator (e.g., Section V-B). Note that the usual graph signal interpolation problem from [6], [8]–[11], [11], [21], [31] can be recovered by taking $\boldsymbol{H} = \boldsymbol{I}$. Then, let the estimate of $\boldsymbol{z}$ be

$$\hat{\boldsymbol{z}} = \boldsymbol{L}\boldsymbol{y}_{\mathcal{S}} = \boldsymbol{L}\boldsymbol{C}\boldsymbol{y}, \tag{7}$$

for some $\boldsymbol{L} \in \mathbb{C}^{N\times|\mathcal{S}|}$, leading to an interpolation error whose covariance matrix is defined as

$$\boldsymbol{K}(\hat{\boldsymbol{z}}) = \mathbb{E}(\boldsymbol{z}-\hat{\boldsymbol{z}})(\boldsymbol{z}-\hat{\boldsymbol{z}})^H. \tag{8}$$

Because $\boldsymbol{L}$ recovers (approximates) $\boldsymbol{z}$ from the samples $\boldsymbol{y}_{\mathcal{S}}$, it is referred to as a *linear interpolation operator* [6], [10], [11].

Given (4)–(8), the optimal graph signal interpolation problem can be cast as

$$\begin{aligned} \underset{\boldsymbol{L}}{\text{minimize}} \quad & \boldsymbol{K}\left(\hat{\boldsymbol{z}}\right) \\ \text{subject to} \quad & \hat{\boldsymbol{z}} = \boldsymbol{L}\boldsymbol{y}_{\mathcal{S}} \end{aligned} \tag{9}$$

where the minimization is with respect to the PSD cone. In other words, if $\boldsymbol{L}^\star$ is a solution of (9) and $\hat{\boldsymbol{z}}^\star = \boldsymbol{L}^\star\boldsymbol{y}_{\mathcal{S}}$, then the covariance matrix in (8) is such that $\boldsymbol{K}(\hat{\boldsymbol{z}}^\star) \preceq \boldsymbol{K}(\hat{\boldsymbol{z}})$ for all $\hat{\boldsymbol{z}}$ as in (7). Note that the problem in (9) is more general than ordinary least-squares estimation since $\text{MSE}(\hat{\boldsymbol{z}}) = \mathbb{E}\left\|\boldsymbol{z}-\hat{\boldsymbol{z}}\right\|_2^2 = \text{Tr}\left[\boldsymbol{K}(\hat{\boldsymbol{z}})\right]$ and $\boldsymbol{K}(\hat{\boldsymbol{z}}^\star) \preceq \boldsymbol{K}(\hat{\boldsymbol{z}}) \Rightarrow \text{MSE}(\hat{\boldsymbol{z}}^\star) \leq \text{MSE}(\hat{\boldsymbol{z}})$ [26].

The following proposition gives the solution of (9) and its interpolation error covariance matrix.

Proposition 1. *Let $\boldsymbol{x} = \boldsymbol{V}_{\mathcal{K}}\bar{\boldsymbol{x}}_{\mathcal{K}}$ be a bandlimited stationary graph signal and $\boldsymbol{y} = \boldsymbol{x}+\boldsymbol{w}$ be its noisy observations, where $\mathbb{E}\,\bar{\boldsymbol{x}}_{\mathcal{K}}\bar{\boldsymbol{x}}_{\mathcal{K}}^H = \boldsymbol{\Lambda}$ and $\mathbb{E}\,\boldsymbol{w}\boldsymbol{w}^H = \boldsymbol{\Lambda}_w$. Then, given a sampling set $\mathcal{S}$, an optimal interpolator $\boldsymbol{L}^\star$ that solves problem* (9) *can be obtained as a solution of*

$$\boldsymbol{L}^\star\boldsymbol{C}\left(\boldsymbol{V}_{\mathcal{K}}\boldsymbol{\Lambda}\boldsymbol{V}_{\mathcal{K}}^H + \boldsymbol{\Lambda}_w\right)\boldsymbol{C}^T = \boldsymbol{H}\boldsymbol{V}_{\mathcal{K}}\boldsymbol{\Lambda}\boldsymbol{V}_{\mathcal{K}}^H\boldsymbol{C}^T. \tag{10}$$

The error covariance matrix of the optimal interpolation $\hat{\boldsymbol{x}}^\star = \boldsymbol{L}^\star\boldsymbol{y}_{\mathcal{S}}$ is given by

$$\boldsymbol{K}^\star(\mathcal{S}) = \boldsymbol{H}\boldsymbol{V}_{\mathcal{K}}\left(\boldsymbol{\Lambda}^{-1} + \sum_{i\in\mathcal{S}}\lambda_{w,i}^{-1}\boldsymbol{v}_i\boldsymbol{v}_i^H\right)^{-1}\boldsymbol{V}_{\mathcal{K}}^H\boldsymbol{H}^H, \tag{11}$$

where $\boldsymbol{v}_i^H$ is the complex conjugate of the i-th row of $\boldsymbol{V}_{\mathcal{K}}$.

Proof. Start by substituting (6) and (7) into the definition of $\boldsymbol{K}$ in (8) to get

$$\boldsymbol{K}(\boldsymbol{L}\boldsymbol{y}_{\mathcal{S}}) = \mathbb{E}(\boldsymbol{H}\boldsymbol{V}_{\mathcal{K}}\bar{\boldsymbol{x}}_{\mathcal{K}} - \boldsymbol{L}\boldsymbol{y}_{\mathcal{S}})(\boldsymbol{H}\boldsymbol{V}_{\mathcal{K}}\bar{\boldsymbol{x}}_{\mathcal{K}} - \boldsymbol{L}\boldsymbol{y}_{\mathcal{S}})^H.$$

Then, using the fact that $\boldsymbol{x}$ is bandlimited and stationary [see (3)], $\boldsymbol{K}$ expands to

$$\begin{aligned} \boldsymbol{K}(\boldsymbol{L}\boldsymbol{y}_{\mathcal{S}}) = {} & \boldsymbol{H}\boldsymbol{V}_{\mathcal{K}}\boldsymbol{\Lambda}\boldsymbol{V}_{\mathcal{K}}^H\boldsymbol{H}^H - \boldsymbol{L}\boldsymbol{C}\boldsymbol{V}_{\mathcal{K}}\boldsymbol{\Lambda}\boldsymbol{V}_{\mathcal{K}}^H\boldsymbol{H}^H \\ & - \boldsymbol{H}\boldsymbol{V}_{\mathcal{K}}\boldsymbol{\Lambda}\boldsymbol{V}_{\mathcal{K}}^H\boldsymbol{C}^T\boldsymbol{L}^H \\ & + \boldsymbol{L}\boldsymbol{C}(\boldsymbol{V}_{\mathcal{K}}\boldsymbol{\Lambda}\boldsymbol{V}_{\mathcal{K}}^H + \boldsymbol{\Lambda}_w)\boldsymbol{C}^T\boldsymbol{L}^H. \end{aligned} \tag{12}$$

From the partial ordering of the PSD cone, $\boldsymbol{L}^\star$ can be obtained by minimizing the scalar cost function

$$J(\boldsymbol{L}) = \boldsymbol{b}^H\boldsymbol{K}(\boldsymbol{L}\boldsymbol{y}_{\mathcal{S}})\boldsymbol{b} \tag{13}$$

simultaneously for all $\boldsymbol{b} \in \mathbb{C}^n$ [26]. Substituting (12) into (13) and setting its derivative with respect to $\boldsymbol{b}^H\boldsymbol{L}$ to zero gives

$$\begin{aligned} &\frac{\partial J(\boldsymbol{L})}{\partial \boldsymbol{b}^H\boldsymbol{L}} = \boldsymbol{0} \Leftrightarrow \\ &\quad \boldsymbol{C}\left(\boldsymbol{V}_{\mathcal{K}}\boldsymbol{\Lambda}\boldsymbol{V}_{\mathcal{K}}^H + \boldsymbol{\Lambda}_w\right)\boldsymbol{C}^T\boldsymbol{L}^H\boldsymbol{b} = \boldsymbol{C}\boldsymbol{V}_{\mathcal{K}}\boldsymbol{\Lambda}\boldsymbol{V}_{\mathcal{K}}^H\boldsymbol{H}^H\boldsymbol{b}. \end{aligned}$$

Since it must hold for all $\boldsymbol{b}$ simultaneously, we obtain (10).

To determine the error covariance matrix $\boldsymbol{K}^\star$ of the optimal interpolator, replace any $\boldsymbol{L}^\star$ satisfying (10) into (12) to get

$$\begin{aligned} \boldsymbol{K}^\star(\boldsymbol{L}^\star\boldsymbol{y}_{\mathcal{S}}) = \boldsymbol{H}\boldsymbol{V}_{\mathcal{K}}\Big\{&\boldsymbol{\Lambda} - \boldsymbol{\Lambda}\boldsymbol{V}_{\mathcal{K}}^H\boldsymbol{C}^T \times \\ &\left[\boldsymbol{C}\left(\boldsymbol{V}_{\mathcal{K}}\boldsymbol{\Lambda}\boldsymbol{V}_{\mathcal{K}}^H + \boldsymbol{\Lambda}_w\right)\boldsymbol{C}^T\right]^{-1}\boldsymbol{C}\boldsymbol{V}_{\mathcal{K}}\boldsymbol{\Lambda}\Big\}\boldsymbol{V}_{\mathcal{K}}^H\boldsymbol{H}^H. \end{aligned} \tag{14}$$

Note that (14) does not depend on $\boldsymbol{L}^\star$ or $\boldsymbol{y}_{\mathcal{S}}$, only on the sampling set $\mathcal{S}$ through the selection matrix $\boldsymbol{C}$. Moreover, since $\boldsymbol{\Lambda}_w$ is diagonal and full rank, $(\boldsymbol{C}\boldsymbol{\Lambda}_w\boldsymbol{C}^T)^{-1} = \boldsymbol{C}\boldsymbol{\Lambda}_w^{-1}\boldsymbol{C}^T$, so that the inverse in (14) always exists. Therefore, using the matrix inversion lemma [27] gives

$$\boldsymbol{K}^\star(\mathcal{S}) = \boldsymbol{H}\boldsymbol{V}_{\mathcal{K}}\left(\boldsymbol{\Lambda}^{-1} + \boldsymbol{V}_{\mathcal{K}}^H\boldsymbol{C}^T\boldsymbol{C}\boldsymbol{\Lambda}_w^{-1}\boldsymbol{C}^T\boldsymbol{C}\boldsymbol{V}_{\mathcal{K}}\right)^{-1}\boldsymbol{V}_{\mathcal{K}}^H\boldsymbol{H}^H. \tag{15}$$

Since $\boldsymbol{C}^T\boldsymbol{C}$ is a diagonal matrix with ones on the indices in $\mathcal{S}$ and zeros everywhere else, (11) is obtained by noting that

$$\boldsymbol{V}_{\mathcal{K}}^H\boldsymbol{C}^T\boldsymbol{C}\boldsymbol{\Lambda}_w^{-1}\boldsymbol{C}^T\boldsymbol{C}\boldsymbol{V}_{\mathcal{K}} = \sum_{i\in\mathcal{S}}\lambda_{w,i}^{-1}\boldsymbol{v}_i\boldsymbol{v}_i^H,$$

for $\boldsymbol{V}_{\mathcal{K}} = \left[\begin{array}{ccc}\boldsymbol{v}_1 & \cdots & \boldsymbol{v}_N\end{array}\right]^H$. ∎

Given a sampling set $\mathcal{S}$, Proposition 1 can be used to optimally estimate a graph signal (or its linear transformation) from its samples. Note that although we assumed that $\boldsymbol{\Lambda}_w$ was full-rank, (10) also holds in the noiseless case ($\boldsymbol{\Lambda}_w = \boldsymbol{0}$), though its solution may no longer be unique. In particular, this happens if the sampling set is not sufficient to determine $\boldsymbol{z}$, i.e., if $\boldsymbol{C}\boldsymbol{V}_{\mathcal{K}}$ is rank-deficient [8]–[11]. In contrast, since $\boldsymbol{\Lambda}_w \succ \boldsymbol{0}$, the matrix on the left-hand side of (10) is always invertible and $\boldsymbol{L}^\star$ is unique for all $\mathcal{S}$. This is similar to the well-known regularization effect of noise in Kalman filtering [26]. The interpolation performance given in (11), however, is not the same for all sampling sets.

B. Sampling Set Selection

Proposition 1 allows us to evaluate the optimal interpolator $\boldsymbol{L}^\star$ that minimizes the interpolation error covariance matrix for a given sampling set. This does not guarantee, however, that there is no other sampling set of the same size for which the interpolation error is smaller. To address this issue, we investigate the *sampling set selection* problem which sets out to find the sampling set that minimizes the interpolation error over all sampling sets. Explicitly, we wish to solve

$$\begin{aligned} \underset{\mathcal{S}\subseteq\mathcal{V}}{\text{minimize}} \quad & \text{MSE}(\mathcal{S}) \\ \text{subject to} \quad & |\mathcal{S}| \leq k \end{aligned} \tag{16}$$

where $\text{MSE}(\mathcal{S}) = \text{Tr}[\boldsymbol{K}^\star(\mathcal{S})]$. Note that in contrast to (9), (16) is a scalar optimization problem. Indeed, optimizing $\boldsymbol{K}^\star$ directly may yield several *minimal* solution. The fact that (9) has a minimum in the PSD cone is a characteristic of stochastic estimation problems that does not hold in general [26].

Another important fact to note is that increasing $\mathcal{S}$ always decreases MSE. This has two important consequences. First, the unconstrained version of (16) is trivial, i.e., its solution is $\mathcal{S} = \mathcal{V}$. Second, it implies that the constraint in (16) is tight, i.e., it can be replaced by the equality constraint $|\mathcal{S}| = k$ without changing the problem solution. This property is a

direct corollary of the following lemma and the monotonicity of the trace operator [32]:

Lemma 1. *The matrix-valued set function $\boldsymbol{K}^\star(\mathcal{S})$ in (11) is monotonically decreasing with respect to the PSD cone, i.e., $\boldsymbol{K}^\star(\mathcal{A}) \succeq \boldsymbol{K}^\star(\mathcal{B})$ whenever $\mathcal{A} \subseteq \mathcal{B}$.*

Proof. Start by noting that $\boldsymbol{K}^\star$ in (11) can be written as

$$\boldsymbol{K}^\star(\mathcal{S}) = \boldsymbol{H}\boldsymbol{V}_{\mathcal{K}}\bar{\boldsymbol{K}}(\mathcal{S})\boldsymbol{V}_{\mathcal{K}}^H\boldsymbol{H}^H,$$

with $\bar{\boldsymbol{K}}(\mathcal{S}) = \left[\boldsymbol{\Lambda}^{-1} + \boldsymbol{R}(\mathcal{S})\right]^{-1}$ and

$$\boldsymbol{R}(\mathcal{S}) = \sum_{i\in\mathcal{S}} \lambda_{w,i}^{-1}\boldsymbol{v}_i\boldsymbol{v}_i^H. \tag{17}$$

Since $\boldsymbol{K}^\star$ and $\bar{\boldsymbol{K}}$ are congruent, suffices to show that $\bar{\boldsymbol{K}}$ is a monotonically decreasing set function [27].

To do so, note that $\bar{\boldsymbol{K}}$ only depends on $\mathcal{S}$ through $\boldsymbol{R}(\mathcal{S})$ and that $\boldsymbol{R}$ is additive, i.e., $\boldsymbol{R}(\mathcal{A}\cup\mathcal{B}) = \boldsymbol{R}(\mathcal{A}) + \boldsymbol{R}(\mathcal{B})$. Then, since $\lambda_{w,i} > 0$, $\boldsymbol{R}$ is a sum of PSD matrices, which implies that $\mathcal{A} \subseteq \mathcal{B} \Rightarrow \boldsymbol{R}(\mathcal{A}) \preceq \boldsymbol{R}(\mathcal{B})$, i.e., $\boldsymbol{R}$ is monotonically increasing. From the antitonicity of the matrix inverse [32], it follows that $\bar{\boldsymbol{K}}$ is monotonically decreasing. ■

Although Lemma 1 reduced the searching space to sampling sets of size k, (16) remains a combinatorial optimization problem: it requires $\binom{n}{k}$ sampling sets to be checked, which is impractical even for moderately small n. In fact, due to the irregularity of the domain of graph signals, sampling set selection is NP-hard in general. It is straightforward to see that it is equivalent to the sensor placement or forward regression problems in [13]–[16].

In the following sections, we address this issue in two ways. First, we derive universal performance bounds that hold for all sampling sets (Section III). These bounds can therefore be used to evaluate the quality of a sampling set or selection heuristics *a posteriori*. Second, we study the greedy sampling algorithm and provide near-optimal guarantees based on the concept of *approximate submodularity* (Section IV). Special cases of these results that considered real homeoscedastic signal and noise ($\boldsymbol{\Lambda} = \sigma_x^2\boldsymbol{I}$ and $\boldsymbol{\Lambda}_w = \sigma_w^2\boldsymbol{I}$) and no transformation of the graph signal ($\boldsymbol{H} = \boldsymbol{I}$) can be found in [21], [31].

III. Universal Bounds on Interpolation MSE

In this section, we derive interpolation performance bounds that hold for all $\mathcal{S}$. These universal bounds can be used to inform the sampling set selection by (i) describing how different factors influence the reconstruction performance and (ii) gauging the quality of sampling set instances. The main result of this section is presented below.

Theorem 1. *Let $\boldsymbol{x} = \boldsymbol{V}_{\mathcal{K}}\bar{\boldsymbol{x}}_{\mathcal{K}}$ be a bandlimited stationary graph signal, $\boldsymbol{y} = \boldsymbol{x} + \boldsymbol{w}$ be its noisy observations, and $\hat{\boldsymbol{z}}^\star = \boldsymbol{L}^\star\boldsymbol{y}_{\mathcal{S}}$ be the minimum MSE interpolation of $\boldsymbol{z} = \boldsymbol{H}\boldsymbol{x}$ based on a sampling set $\mathcal{S}$. Also, take $\mathbb{E}\,\bar{\boldsymbol{x}}_{\mathcal{K}}\bar{\boldsymbol{x}}_{\mathcal{K}}^H = \boldsymbol{\Lambda}$, $\mathbb{E}\,\boldsymbol{w}\boldsymbol{w}^H = \boldsymbol{\Lambda}_w$, and $\boldsymbol{W} = \boldsymbol{V}_{\mathcal{K}}^H\boldsymbol{H}^H\boldsymbol{H}\boldsymbol{V}_{\mathcal{K}}$. Then, if $\boldsymbol{W} \succ 0$, the reconstruction error $\mathrm{MSE}(\mathcal{S}) = \mathbb{E}\left\|\boldsymbol{z} - \hat{\boldsymbol{z}}^\star\right\|^2 = \mathrm{Tr}[\boldsymbol{K}^\star(\mathcal{S})]$ is bounded by*

$$\frac{|\mathcal{K}|^2}{\mathrm{Tr}\left[(\boldsymbol{W}\boldsymbol{\Lambda})^{-1}\right] + \bar{\ell}_{|\mathcal{S}|}} \leq \mathrm{MSE}(\hat{\boldsymbol{z}}^\star) \leq \mathrm{Tr}\left(\boldsymbol{W}\boldsymbol{\Lambda}\right), \tag{18}$$

where $\bar{\ell}_m$ is the sum of the m largest weighted structural SNRs $\ell_i = \lambda_{w,i}^{-1}\left\|\boldsymbol{v}_i\right\|_{\boldsymbol{W}^{-1}}^2$, with $\boldsymbol{v}_i^H$ the i-th row of $\boldsymbol{V}_{\mathcal{K}}$ and $\|\boldsymbol{x}\|_{\boldsymbol{A}}^2 = \boldsymbol{x}^H\boldsymbol{A}\boldsymbol{x}$. Explicitly, $\bar{\ell}_m = \max_{\mathcal{X}:|\mathcal{X}|=m}\sum_{j\in\mathcal{X}}\ell_j$.

Proof. Start with the upper bound that is achieved for an empty sampling set, i.e., for $\mathcal{S} = \{\}$. Indeed, recall from Lemma 1 that $\boldsymbol{K}^\star$ is a monotone decreasing set function, i.e., it achieves its maximum for the empty set. Thus, it holds that $\boldsymbol{K}(\hat{\boldsymbol{x}}^\star) \preceq \boldsymbol{H}\boldsymbol{V}_{\mathcal{K}}\boldsymbol{\Lambda}\boldsymbol{V}_{\mathcal{K}}^H\boldsymbol{H}^H$, from which the upper bound in (18) follows by the monotonicity of the trace operator [32].

To obtain the lower bound, start by using (11) to get

$$\mathrm{MSE}(\mathcal{S}) = \mathrm{Tr}\left[\boldsymbol{W}\left(\boldsymbol{\Lambda}^{-1} + \sum_{i\in\mathcal{S}}\lambda_{w,i}^{-1}\boldsymbol{v}_i\boldsymbol{v}_i^H\right)^{-1}\right], \tag{19}$$

where we used the circular commutation property of the trace. Then, since the trace of a matrix is the sum of its eigenvalues, the arithmetic/harmonic means inequality can be used to get, for any $N \times N$ positive-definite matrix $\boldsymbol{X}$,

$$\mathrm{Tr}\left(\boldsymbol{X}\right) \geq \frac{N^2}{\mathrm{Tr}\left(\boldsymbol{X}^{-1}\right)},$$

with equality if and only if $\boldsymbol{X} = \gamma\boldsymbol{I}$, $\gamma > 0$ [27]. Since $\boldsymbol{W} \succ 0$, the matrix in (19) is positive-definite and on has

$$\mathrm{MSE}(\mathcal{S}) \geq \frac{|\mathcal{K}|^2}{\mathrm{Tr}\left[(\boldsymbol{W}\boldsymbol{\Lambda})^{-1}\right] + \mathrm{Tr}\left[\boldsymbol{W}^{-1}\left(\sum_{i\in\mathcal{S}}\lambda_{w,i}^{-1}\boldsymbol{v}_i\boldsymbol{v}_i^H\right)\right]},$$

which from the commutation property of the trace gives

$$\mathrm{MSE}(\mathcal{S}) \geq \frac{|\mathcal{K}|^2}{\mathrm{Tr}\left[(\boldsymbol{W}\boldsymbol{\Lambda})^{-1}\right] + \sum_{i\in\mathcal{S}}\lambda_{w,i}^{-1}\left\|\boldsymbol{v}_i\right\|_{\boldsymbol{W}^{-1}}^2},$$

where we used the weighted norm $\|\boldsymbol{x}\|_{\boldsymbol{A}}^2 = \boldsymbol{x}^H\boldsymbol{A}\boldsymbol{x}$. Finally, replacing the sum in the denominator by its maximum possible value $\bar{\ell}_{|\mathcal{S}|}$ gives the desired lower bound in (18). ■

The bounds in (18) were derived by taking the graph signal to be stochastic, so that the expectation in the MSE is taken over realizations of the signal and the bounds holds for all sampling sets $\mathcal{S} \subseteq \mathcal{V}$. Also, it is worth noting that (18) depends only on statistics of the graph signal ($\boldsymbol{\Lambda}$, $\boldsymbol{\Lambda}_w$, and $\mathcal{K}$), graph signal transform ($\boldsymbol{H}$), structural properties of the underlying graph ($\boldsymbol{V}$), and the sampling set size ($|\mathcal{S}|$). These are all quantities known *a priori*, i.e., before the sampling occurs.

As expected, (18) decreases with the sampling set size. The rate of decay, however, depends on the *weighted structural SNRs* $\{\ell_i\}$. These quantities represent the relation between the signal of interest and the noise at each node, taking into account the structure of the graph and the subspace of interest ($\mathrm{colspan}(\boldsymbol{V}_{\mathcal{K}})$). Moreover, they are related to statistical estimates such as the leverage score and the Mahalanobis distance in regression. In a sequential sampling scheme, their value can be used to inform whether a new sample is worth acquiring by bounding the possible MSE improvement. A bound on the decay rate can also be obtained using the fact that $\bar{\ell}_m \leq m\ell_{\max}$ for $\ell_{\max} = \max_i \ell_i$. Then, if the sampling set is chosen so as to uniquely determine the graph signal, i.e., $|\mathcal{S}| \geq |\mathcal{K}|$, (18) reduces to

$$\mathrm{MSE}(\mathcal{S}) \geq \frac{|\mathcal{K}|}{\lambda_{\min}(\boldsymbol{W}\boldsymbol{\Lambda})^{-1} + \ell_{\max}}, \tag{20}$$

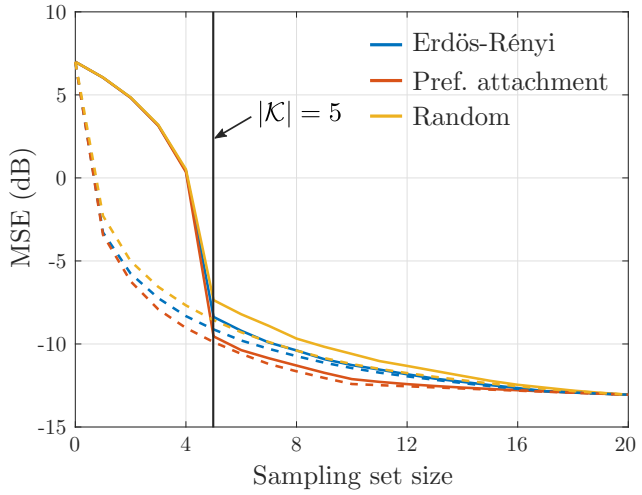

Figure 1. Comparison between (18) (dashed lines) and minimum MSE (solid lines) for reconstructing graph signals ($\boldsymbol{H} = \boldsymbol{I}$) on random graphs ($n = 20$)

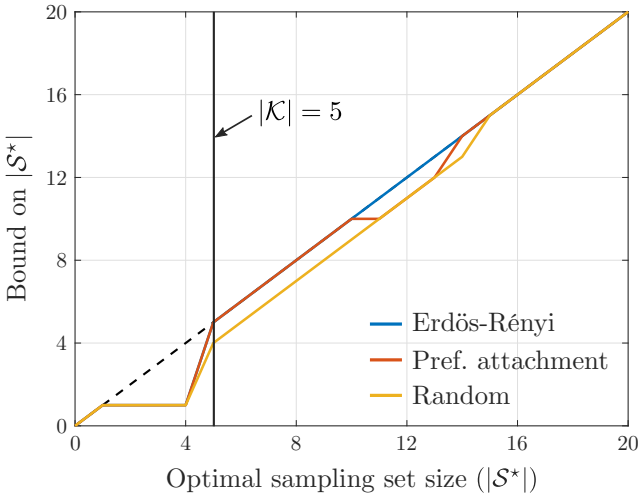

Figure 2. Comparison between (21) (dashed lines) and optimal sampling set size (solid lines) for reconstructing graph signals ($\boldsymbol{H} = \boldsymbol{I}$) on random graphs ($n = 20$)

where the fact that $\mathrm{Tr}\left[(\boldsymbol{W}\boldsymbol{\Lambda})^{-1}\right] \leq |\mathcal{K}|\lambda_{\min}(\boldsymbol{W}\boldsymbol{\Lambda})^{-1}$ was used. It is clear from (20) that the reconstruction error increases linearly with the bandwidth of the graph signal, which is a fundamental limitation for large dimensional signals. It also shows the importance of working with low bandwidth signals and, consequently, of appropriately identifying the signal's underlying graph.

Although these observations give insights into graph signal interpolation, one of the main motivation behind Theorem 1 is addressing the issue of sampling set selection. Towards this end, we propose the following corollary:

Corollary 1. *For any graph signal and its interpolation as in Theorem 1, any sampling set $\mathcal{S}$ for which* $\mathrm{MSE}(\mathcal{S}) \leq \eta$ *satisfies*

$$\bar{\ell}_{|\mathcal{S}|} \geq \frac{|\mathcal{K}|^2 - \eta\,\mathrm{Tr}\left[(\boldsymbol{W}\boldsymbol{\Lambda})^{-1}\right]}{\eta}. \tag{21}$$

Since $\bar{\ell}_{|\mathcal{S}|} \leq |\mathcal{S}|\ell_{max}$, it also holds that

$$|\mathcal{S}| \geq \frac{|\mathcal{K}|^2 - \eta\,\mathrm{Tr}\left[(\boldsymbol{W}\boldsymbol{\Lambda})^{-1}\right]}{\eta\ell_{max}}. \tag{22}$$

Corollary 1 gives a lower bound on the number of samples needed to achieve a desired MSE. From (22), note that the minimum number of samples increases as the MSE decreases. Moreover, although (22) suggest that the sample set size required to achieve a certain MSE grows with $\mathcal{O}(|\mathcal{K}|^2)$, it is not necessarily the case. Indeed, recall that $\ell_{\max}$ is a function of $|\mathcal{K}|$ (through $\|\boldsymbol{v}_i\|$ and $\boldsymbol{V}_{\mathcal{K}}$). Still, as in the noiseless case, the signal bandwidth is a dominating factor in the determination of the minimum sampling set size.

Although (22) characterizes the overall behavior of the sampling set size, it is not informative in practice because it largely underestimates $|\mathcal{S}|$. On the other hand, (21) yields a tighter bound which can be used, together with (18), to evaluate a sampling set or sampling technique for direct reconstruction of a graph signal ($\boldsymbol{H} = \boldsymbol{I}$). Indeed, Figures 1 and 2 compares (18) and (21) to the minimum interpolation MSE and optimal set size, found by exhaustive search, for three graph

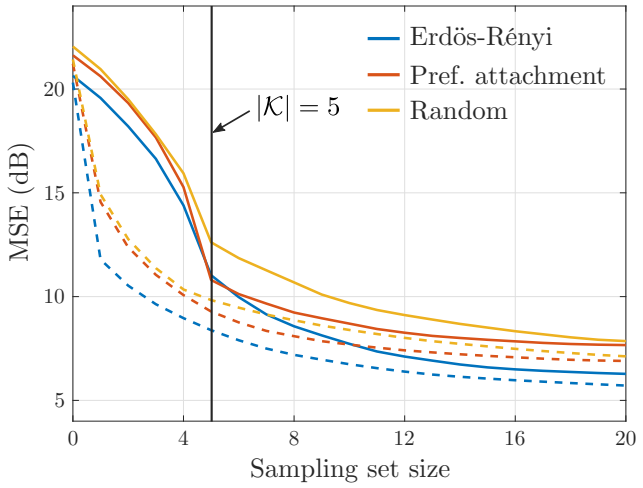

Figure 3. Comparison between (18) (dashed lines) and minimum MSE (solid lines) for arbitrary $\boldsymbol{H}$ ($n = 20$)

models ($n = 20$): Erdös-Rényi, preferential attachment, and a random undirected graph with weights uniformly distributed in $[0, 1]$ (see details of these models in Section V-A). The graph signal is assumed to be homeoscedastic with $\boldsymbol{\Lambda} = \boldsymbol{I}$ and $\boldsymbol{\Lambda}_w = \sigma_w^2\boldsymbol{I}$, $\sigma_w^2 = 10^{-2}$. Note that the bounds are conservative for $|\mathcal{S}| < |\mathcal{K}|$, but become tighter as $|\mathcal{S}|$ increases. This is because the inequality used to derive (18) becomes tighter as the eigenvalues of $\boldsymbol{K}^\star$ become more similar.

When $\boldsymbol{H}$ is arbitrary and variance of signal and noise can vary across nodes, the eigenvalues of $\boldsymbol{K}^\star$ can become different from each another and deteriorate the bound in (18). This is illustrated in Figures 3 and 4, where $\boldsymbol{H}$ was taken as a 30×20 whose entries are zero-mean unit variance Gaussian random variables, $\boldsymbol{\Lambda} = \boldsymbol{I}$, and the noise variance was uniformly distributed in $[10^{-3}, 10^{-1}]$.

Remark 1. Bounds on the interpolation MSE of graph sampling techniques have also been derived in [6], [11]. These works consider randomized sampling set selection schemes,

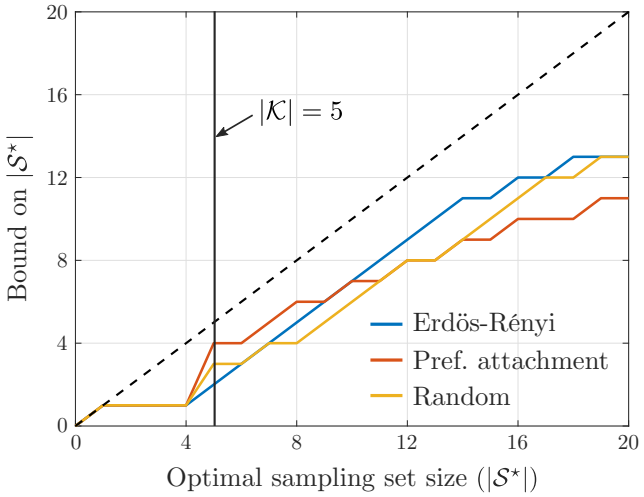

Figure 4. Comparison between (21) (dashed lines) and optimal sampling set size (solid lines) for arbitrary $\boldsymbol{H}$ ($n = 20$)

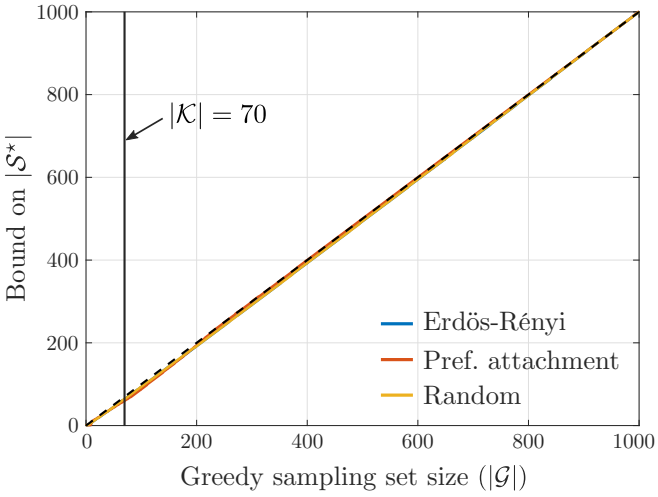

Figure 5. Evaluating sampling sets obtained by greedy sampling (solid lines) and (21) (dashed lines) for different random graphs ($n = 1000$)

including uniform and leverage score sampling, and derive performance bounds on the optimal sampling distributions and interpolation error. The bounds in Theorem 1 and Corollary 1 differ from those in [6], [11] in that the latter take the spectrum of the graph signal to be deterministic and the sampling to be random. Thus, these bounds hold in expectation over different sampling realizations for a specific randomized strategy. The bounds in (18) hold in expectation over realizations of the signal and apply to any sampling strategy.

IV. Near-optimal Sampling Set Selection

Although the bounds from Section III can be used to evaluate specific sampling set instances, they do not provide performance guarantees for any sampling strategies in general. To do that, this section studies a specific sampling scheme, namely *greedy sampling set selection*, and derives near-optimality results that hold for all problem instances.

Greedy sampling set selection is ubiquitous in GSP and has proven successful in many applications [8]–[12], [18]. This is illustrated in Figure 5 which uses the bounds derived in (21) to assess the quality of sampling sets obtained by greedily minimizing the MSE (see Algorithm 2) on larger instances ($n = 1000$) of the three random graph models found in Figures 1 to 4. Note that the final greedy sampling set size remains within 10% of the lower bound in these realizations.

Despite strong empirical evidences, typical performance guarantees for greedy search do not hold for greedy sampling set selection. Indeed, the well-established result from [33] states that greedy minimization (Algorithm 1) yields guaranteed near-optimal results for monotonically decreasing and supermodular set functions. The MSE, however, is not supermodular in general. This can be seen from [15, Thm. 2.4] and the fact that $f(t) = t^{-2}$ is not operator antitone [32]. Thus, although greedily minimizing the MSE appears to work in practice, there has yet to be a theoretical justification for it. The following sections bridge this gap by expanding the notion of approximate supermodularity introduced in [21] and updating the near-optimality bound from [33] to account for this class of functions. This novel framework then allows near-optimality bounds to be derived for the MSE.

Algorithm 1 Greedy minimization

$\mathcal{G}_0 = \{\}$
for $j = 1, \dots, \ell$
  $u = \underset{s \in \mathcal{V} \setminus \mathcal{G}_{j-1}}{\operatorname{argmin}} \; f\left(\mathcal{G}_{j-1} \cup \{s\}\right)$
  $\mathcal{G}_j = \mathcal{G}_{j-1} \cup \{u\}$
end

A. Approximate supermodularity and greedy minimization

Supermodularity (and its dual *submodularity*) encodes the "diminishing returns" property of certain functions that leads to bounds on the suboptimality of their greedy minimization [33]. Well-known supermodular functions include the rank, $\log\det$, or Von-Neumann entropy of a matrix [20]. Still, supermodularity is a stringent condition. In particular, it does not hold for the MSE in (16). To provide new suboptimality bounds for greedy minimization, we therefore define the concept of approximate supermodularity.

A set function $f : 2^{\mathcal{V}} \to \mathbb{R}$ is α-*supermodular* if for all sets $\mathcal{A} \subseteq \mathcal{B} \subseteq \mathcal{V}$ and all $u \notin \mathcal{B}$ it holds that

$$f\left(\mathcal{A} \cup \{u\}\right) - f\left(\mathcal{A}\right) \leq \alpha \left[f\left(\mathcal{B} \cup \{u\}\right) - f\left(\mathcal{B}\right)\right], \qquad (23)$$

for $\alpha \geq 0$. We say f is α-*submodular* if $-f$ is α-supermodular. For $\alpha \geq 1$, (23) is equivalent to the traditional definition of supermodularity, in which case we refer to the function simply as *supermodular*/*submodular* [20]. For $\alpha \in [0, 1)$, however, f is said to be *approximately supermodular*/*submodular*. Notice that (23) always holds for $\alpha = 0$ if f is monotone decreasing. Indeed, $f\left(\mathcal{A} \cup \{u\}\right) - f\left(\mathcal{A}\right) \leq 0$ in this case. Thus, α-supermodularity is only of interest when α takes the largest value for which (23) holds, i.e.,

$$\alpha = \min_{\substack{\mathcal{A} \subseteq \mathcal{B} \subseteq \mathcal{V} \\ v \notin \mathcal{B}}} \frac{f(\mathcal{A} \cup \{v\}) - f(\mathcal{A})}{f(\mathcal{B} \cup \{v\}) - f(\mathcal{B})}. \qquad (24)$$

Before proceeding, it is worth noting that α is related to the *submodularity ratio* from [14]. It is, however, more amenable to give explicit bounds on its value (see Section IV-B). Indeed, the submodularity ratio bounds derived in [14] depend on the minimum sparse eigenvalue of a matrix, which cannot be evaluated efficiently. Due to the relation between α and the submodularity ratio, it is not surprising that similar near-optimal bounds hold for α-supermodular functions.

Theorem 2. *Let $f^\star = f(\mathcal{S}^\star)$ be the optimal value of the problem*

$$\underset{\mathcal{S}\subseteq\mathcal{V},\ |\mathcal{S}|=k}{\text{minimize}} \quad f(\mathcal{S}) \tag{25}$$

and $\mathcal{G}_\ell$ be its greedy solution obtained using Algorithm 1. If f is (i) monotone decreasing and (ii) α-supermodular, then

$$\frac{f(\mathcal{G}_\ell) - f^\star}{f(\{\}) - f^\star} \leq \left(1 - \frac{\alpha}{k}\right)^\ell \leq e^{-\alpha\ell/k}. \tag{26}$$

If f is normalized, i.e., $f(\{\}) = 0$, (26) reduces to

$$\frac{f(\mathcal{G}_\ell)}{f^\star} \leq 1 - e^{-\alpha\ell/k}.$$

Proof. Using the fact that f is monotone decreasing, it holds for every set $\mathcal{G}_j$ that

$$f(\mathcal{S}^\star) \geq f(\mathcal{S}^\star \cup \mathcal{G}_j).$$

Using a telescopic sum then gives

$$f(\mathcal{S}^\star) \geq f(\mathcal{G}_j) + \sum_{i=1}^{k} f(\mathcal{T}_{i-1} \cup \{s_i^\star\}) - f(\mathcal{T}_{i-1}), \tag{27}$$

where $\mathcal{T}_i = \mathcal{G}_j \cup \{s_1^\star, \ldots, s_i^\star\}$ and $s_i^\star$ is the i-th element of $\mathcal{S}^\star$. Since f is α-supermodular and $\mathcal{G}_j \subset \mathcal{T}_i$ for all i, the incremental gains in the summation in (27) can be bounded using (23) to get

$$f(\mathcal{S}^\star) \geq f(\mathcal{G}_j) + \alpha^{-1} \sum_{i=1}^{k} f(\mathcal{G}_j \cup \{s_i^\star\}) - f(\mathcal{G}_j).$$

Finally, given that $\mathcal{G}_{j+1} = \mathcal{G}_j \cup \{u\}$ is chosen to minimize $f(\mathcal{G}_{j+1})$ (see Algorithm 1),

$$f(\mathcal{S}^\star) \geq f(\mathcal{G}_j) + \alpha^{-1} k \left[f(\mathcal{G}_{j+1}) - f(\mathcal{G}_j)\right]. \tag{28}$$

To obtain a recursion, let $\delta_j = f(\mathcal{G}_j) - f(\mathcal{S}^\star)$ so that (28) becomes

$$\delta_j \leq \alpha^{-1} k \left[\delta_j - \delta_{j+1}\right] \Rightarrow \delta_{j+1} \leq \left(1 - \frac{1}{\alpha^{-1}k}\right) \delta_j.$$

Noting that $\delta_0 = f(\{\}) - f(\mathcal{S}^\star)$, we can solve this recursion to get

$$\frac{f(\mathcal{G}_\ell) - f(\mathcal{S}^\star)}{f(\{\}) - f(\mathcal{S}^\star)} \leq \left(1 - \frac{\alpha}{k}\right)^\ell.$$

Using the fact that $1 - x \leq e^{-x}$ yields (26). ∎

Theorem 2 bounds the relative suboptimality of the greedy solution to problem (25) when f is decreasing and α-supermodular. Under these conditions, it guarantees a minimum improvement of the greedy solution over the empty set. What is more, it quantifies the effect of relaxing the supermodularity hypothesis in (23). Indeed, when f is supermodular ($\alpha = 1$) and the greedy search in Algorithm 1 is repeated k times ($\ell = k$), we recover the $e^{-1} \approx 0.37$ guarantee from [33]. On the other hand, if f is not supermodular ($\alpha < 1$), (26) shows that the same 37% guarantee can be obtained by greedily selecting a set of size $\alpha^{-1}k$. Thus, α not only quantifies how much f violates supermodularity, but also gives a factor by which a solution set must increase to maintain supermodular near-optimality. It is worth noting that, as with the original bound in [33], (26) is not tight and that better results are common in practice (see Section V-A).

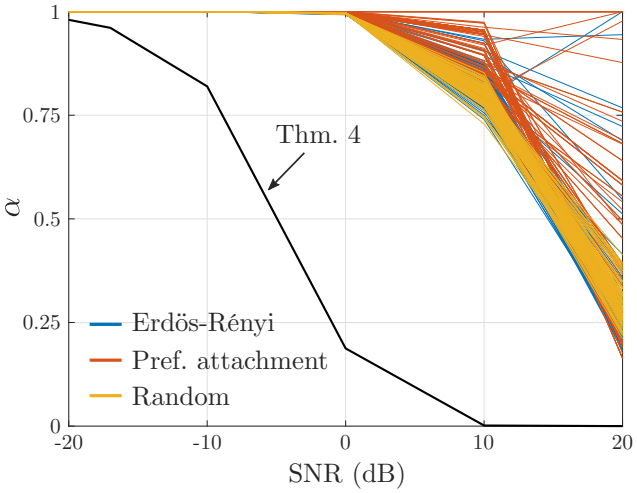

Figure 6. Comparison between the bound in Theorem 4 and α

In the sequel, we show that $\text{MSE}(\mathcal{S})$ is a monotone decreasing and α-supermodular functions of the sampling set $\mathcal{S}$. We also provide an explicit lower bound on α as a function of the SNR. This result simultaneously provide near-optimal performance guarantees based on Theorem 2 and sheds light on why greedy algorithms have been so successful in GSP applications.

B. Near-optimality of greedy sampling set selection

The main result of this section can be stated as:

Theorem 3. *The scalar set functions $\text{MSE}(\mathcal{S}) = \text{Tr}[\boldsymbol{K}^\star(\mathcal{S})]$ is (i) monotone decreasing and (ii) α-supermodular with*

$$\alpha \geq \frac{\lambda_{\max}(\boldsymbol{\Lambda}_w)^{-1} + \mu_{\max}^{-1}}{\lambda_{\max}(\boldsymbol{\Lambda}_w)^{-1} + \mu_{\min}^{-1}} \frac{\mu_{\min}^2}{\kappa_2(\boldsymbol{W})\,\mu_{\max}^2},$$

where $\mu_{\min} \leq \lambda_{\min}\left[\boldsymbol{\Lambda}^{-1}\right]$, $\mu_{\max} \geq \lambda_{\max}\left[\boldsymbol{\Lambda}^{-1} + \boldsymbol{V}_{\mathcal{K}}^H \boldsymbol{\Lambda}_w^{-1} \boldsymbol{V}_{\mathcal{K}}\right]$, and $\kappa_2(\boldsymbol{W})$ is the 2-norm condition number of $\boldsymbol{W}$ [27].

The proof of Theorem 3 is deferred to Appendix A. For clarity, we instead study the case in which both signal and noise are homoscedastic, i.e., $\boldsymbol{\Lambda} = \sigma_x^2 \boldsymbol{I}$ and $\boldsymbol{\Lambda}_w = \sigma_w^2 \boldsymbol{I}$. Although both proofs follow similar steps, the result in the latter case is more straightforward to understand and interpret. Explicitly, this section focuses on the following result:

Theorem 4. *Assuming $\boldsymbol{\Lambda} = \sigma_x^2 \boldsymbol{I}$ and $\boldsymbol{\Lambda}_w = \sigma_w^2 \boldsymbol{I}$, the set functions $\text{MSE}(\mathcal{S}) = \text{Tr}[\boldsymbol{K}^\star(\mathcal{S})]$ is (i) monotone decreasing and (ii) α-supermodular with*

$$\alpha \geq \frac{1+2\gamma}{\kappa_2(\boldsymbol{W})\,(1+\gamma)^4}, \quad \textit{for} \quad \gamma = \frac{\sigma_x^2}{\sigma_w^2}, \tag{29}$$

where $\kappa_2(\boldsymbol{W})$ is the 2-norm condition number of $\boldsymbol{W}$ [27].

Theorems 3 and 4 establish that a near-optimal solution to the sampling set selection problem in (16) can be obtained efficiently using greedy search. Though strong empirical evidence exist that greedily minimizing the MSE yields good results in contexts such as regression, dictionary learning, and graph signal processing [8]–[11], [14], [18], this result is counter-intuitive given that the MSE is not supermodular in general. For instance, restrictive and often unrealistic conditions on data distribution are required to obtain supermodularity in the context of regression [14].

Theorems 3 and 4 reconcile the empirical success of greedy sampling set selection and the non-supermodularity of the MSE by bounding the suboptimality of greedy sampling. Theorem 4, in particular, gives a simple bound on α in terms of the SNR and the condition number of $\boldsymbol{W}$ that gives clear insights into its behavior. Indeed, as $\gamma \to \infty$ and we approach the noiseless case, $\alpha \to 0$. This is expected as in the noiseless case almost every set of size $|\mathcal{K}|$ achieves perfect reconstruction, so that the choice of sampling nodes is irrelevant. On the other hand, $\alpha \to 1$ as $\gamma \to 0$, i.e., the MSE becomes closer to supermodular as the SNR decreases. Given that reconstruction errors are small for high SNR, Theorem 2 guarantees that greedy sampling performs well when it is most needed. Similar trends can be observed in the more general setting of Theorem 3. These observations are illustrated in Fig. 6 that compares the bound in Theorem 4 to the true value of α for the MSE (found by exhaustive search) in 100 realizations of random graphs (see Section V-A for details).

We can now proceed with the proof of Theorem 4 after stating a pertinent remark.

Remark 2. Since the MSE is *not* submodular, it is common to see surrogate submodular figures of merit used instead, specially in statistics and experiment design [13]–[15], [20]. In particular, the log-determinant $\log\det[\boldsymbol{K}^\star(\mathcal{S})]$ is a common alternative to the objective $\text{MSE}(\mathcal{S}) = \text{Tr}[\boldsymbol{K}^\star(\mathcal{S})]$ used in (16). This is justified because the $\log\det[\boldsymbol{K}^\star(\mathcal{S})]$ is proportional to the volume of the confidence ellipsoids of the estimate when the data is Gaussian [15], [34]. This choice of objective is also common in the sensor placement literature due to its relation to information theoretic measures, such as entropy and mutual information [13]. By replacing the trace operator in (16) by the $\log\det$, the problem becomes a supermodular function minimization that can be efficiently approximated using greedy search, as shown in [12], [21]. We remark that minimizing the $\log\det$ of the error covariance matrix and the MSE are not equivalent problems.

C. Proof of Theorem 4

Start by noticing that part (i) stems directly from Lemma 1. Indeed, the monotonicity of the trace implies that $\boldsymbol{X} \succeq \boldsymbol{Y} \Rightarrow \text{Tr}(\boldsymbol{X}) \geq \text{Tr}(\boldsymbol{Y})$, for any PSD matrices $\boldsymbol{X}$ and $\boldsymbol{Y}$.

Then, to obtain part (ii), use the homeoscedasticity assumption to rewrite (11) as

$$\boldsymbol{K}^\star(\mathcal{S}) = \sigma_x^2 \boldsymbol{H} \boldsymbol{V}_{\mathcal{K}} \boldsymbol{Z}(\mathcal{S})^{-1} \boldsymbol{V}_{\mathcal{K}}^H \boldsymbol{H}^H, \tag{30}$$

where $\boldsymbol{Z}(\mathcal{S}) = \boldsymbol{I} + \gamma \sum_{i\in\mathcal{S}} \boldsymbol{v}_i \boldsymbol{v}_i^H$ and $\gamma = \sigma_x^2/\sigma_w^2$ is the SNR. Then, proceed to obtain a closed form expression for the increments in (24) by using (30) to get

$$f(\mathcal{A}\cup\{u\}) - f(\mathcal{A}) = \sigma_x^2 \,\text{Tr}\left[\boldsymbol{W}\left(\boldsymbol{Z}(\mathcal{A}) + \gamma \boldsymbol{v}_u \boldsymbol{v}_u^H\right)^{-1} - \boldsymbol{W}\boldsymbol{Z}(\mathcal{A})^{-1}\right].$$

From the matrix inversion lemma [27], this expression reduces to

$$f(\mathcal{A}\cup\{u\}) - f(\mathcal{A}) = -\sigma_x^2 \,\text{Tr}\left[\boldsymbol{W}\frac{\boldsymbol{Z}(\mathcal{A})^{-1}\boldsymbol{v}_u \boldsymbol{v}_u^H \boldsymbol{Z}(\mathcal{A})^{-1}}{\gamma^{-1} + \boldsymbol{v}_u^H \boldsymbol{Z}(\mathcal{A})^{-1}\boldsymbol{v}_u}\right],$$

which using the commutation property of the trace yields

$$f(\mathcal{A}\cup\{u\}) - f(\mathcal{A}) = -\sigma_x^2 \frac{\boldsymbol{v}_u^H \boldsymbol{Z}(\mathcal{A})^{-1}\boldsymbol{W}\boldsymbol{Z}(\mathcal{A})^{-1}\boldsymbol{v}_u}{\gamma^{-1} + \boldsymbol{v}_u^H \boldsymbol{Z}(\mathcal{A})^{-1}\boldsymbol{v}_u}. \tag{31}$$

From (31), the expression for α in (24) becomes

$$\alpha = \min_{\substack{\mathcal{A}\subseteq\mathcal{B}\subset\mathcal{V}\\ v\notin\mathcal{B}}} \frac{\gamma^{-1} + \boldsymbol{v}_u^H \boldsymbol{Z}(\mathcal{B})^{-1}\boldsymbol{v}_u}{\gamma^{-1} + \boldsymbol{v}_u^H \boldsymbol{Z}(\mathcal{A})^{-1}\boldsymbol{v}_u} \frac{\boldsymbol{v}_u^H \boldsymbol{Z}(\mathcal{A})^{-1}\boldsymbol{W}\boldsymbol{Z}(\mathcal{A})^{-1}\boldsymbol{v}_u}{\boldsymbol{v}_u^H \boldsymbol{Z}(\mathcal{B})^{-1}\boldsymbol{W}\boldsymbol{Z}(\mathcal{B})^{-1}\boldsymbol{v}_u}. \tag{32}$$

To bound (32), first notice that for any set $\mathcal{X} \subseteq \mathcal{V}$

$$1 \leq \lambda_{\min}\left[\boldsymbol{Z}(\mathcal{X})\right] \leq \lambda_{\max}\left[\boldsymbol{Z}(\mathcal{X})\right] \leq 1+\gamma, \tag{33}$$

where $\lambda_{\min}$ and $\lambda_{\max}$ denote the minimum and maximum eigenvalues of a matrix. These bounds are achieved for the empty set and $\mathcal{V}$, respectively. Then, using the Rayleigh quotient inequalities [27]

$$\|\boldsymbol{b}\|_2^2 \,\lambda_{\min}(\boldsymbol{A}) \leq \boldsymbol{b}^H \boldsymbol{A}\boldsymbol{b} \leq \|\boldsymbol{b}\|_2^2 \,\lambda_{\max}(\boldsymbol{A}),$$

we get that (32) is bounded by

$$\alpha \geq \frac{\gamma^{-1} + \|\boldsymbol{v}_u\|_2^2\,(1+\gamma)^{-1}}{\gamma^{-1} + \|\boldsymbol{v}_u\|_2^2} \cdot \frac{\lambda_{\min}\left[\boldsymbol{Z}(\mathcal{A})^{-1}\boldsymbol{W}\boldsymbol{Z}(\mathcal{A})^{-1}\right]}{\lambda_{\max}\left[\boldsymbol{Z}(\mathcal{B})^{-1}\boldsymbol{W}\boldsymbol{Z}(\mathcal{B})^{-1}\right]}.$$

Using the fact that $\lambda_{\max}(\boldsymbol{A}\boldsymbol{B}) \leq \lambda_{\max}(\boldsymbol{A})\lambda_{\max}(\boldsymbol{B})$ and $\lambda_{\min}(\boldsymbol{A}\boldsymbol{B}) \geq \lambda_{\min}(\boldsymbol{A})\lambda_{\min}(\boldsymbol{B})$ [35], this expression simplifies to

$$\alpha \geq \frac{\gamma^{-1} + 1 + \|\boldsymbol{v}_u\|_2^2}{\gamma^{-1} + \|\boldsymbol{v}_u\|_2^2} \cdot \frac{(1+\gamma)^{-3}}{\kappa_2(\boldsymbol{W})} \triangleq \alpha'. \tag{34}$$

where $\kappa_2(\boldsymbol{W}) = \lambda_{\max}(\boldsymbol{W})/\lambda_{\min}(\boldsymbol{W})$ is the 2-norm condition number of $\boldsymbol{W}$ [27].

Finally, to obtain the expression in Theorem 4, notice that (34) is decreasing with respect to $\|\boldsymbol{v}_u\|_2^2$. Indeed, since $\kappa_2 \geq 1$ and $\gamma \geq 0$,

$$\frac{\partial \alpha'}{\partial \|\boldsymbol{v}_u\|_2^2} = \frac{-(1+\gamma)^{-3}}{\kappa_2(\boldsymbol{W})\left(\gamma^{-1} + \|\boldsymbol{v}_u\|_2^2\right)^2} \leq 0.$$

Given that $\boldsymbol{v}_u^H$ is a row of $\boldsymbol{V}_{\mathcal{K}}$, i.e., it is composed of a subset of elements from a unit vector, $\|\boldsymbol{v}_u\|_2^2 \leq 1$ and we obtain the result in (29). ∎

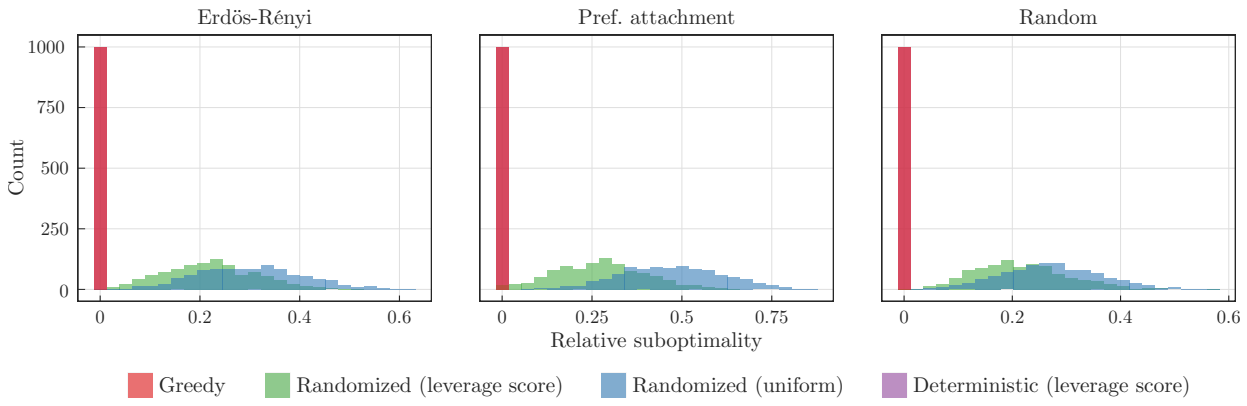

Figure 7. Relative suboptimality of sampling schemes for low SNR (SNR $= -20$ dB)

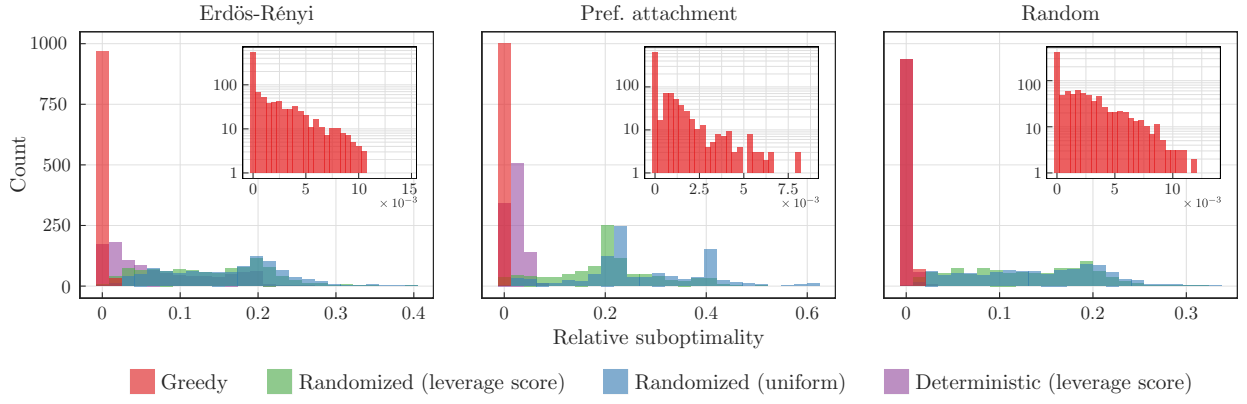

Figure 8. Relative suboptimality of sampling schemes for high SNR (SNR $= 20$ dB)

V. Numerical Examples and Applications

Before proceeding with the simulations, the complexity issue of greedy sampling set selection needs to be addressed. The greedy search in Algorithm 1 requires $\mathcal{O}(n\ell c_f)$ operations, where c_f is the cost of evaluating the cost function f. As it is, problem (16) has $c_f = \mathcal{O}(|\mathcal{K}|^3)$. It can, however, be reduced using the matrix inversion lemma [27].

Indeed, start by noticing that the first step of the greedy approximation of problem (16) involves finding (see Algorithm 1)

$$u = \operatorname*{argmin}_{s\in\mathcal{V}} \operatorname{Tr}\left[\boldsymbol{K}^\star\left(\mathcal{G}_{j-1}\cup\{s\}\right)\right],$$

which, using the definition of $\boldsymbol{K}^\star$ in (11) and the circular commutation property of the trace, requires the evaluation of

$$\operatorname{Tr}\left[\boldsymbol{K}^\star\left(\mathcal{G}_{j-1}\cup\{s\}\right)\right] = \operatorname{Tr}\left[\boldsymbol{W}\left(\boldsymbol{\Lambda}^{-1} + \sum_{i\in\mathcal{G}_{j-1}} \lambda_{w,i}^{-1}\boldsymbol{v}_i\boldsymbol{v}_i^H + \lambda_{w,s}^{-1}\boldsymbol{v}_s\boldsymbol{v}_s^H\right)^{-1}\right],$$

where once again $\boldsymbol{W} = \boldsymbol{V}_{\mathcal{K}}^H\boldsymbol{H}^H\boldsymbol{H}\boldsymbol{V}_{\mathcal{K}}$. Letting $\boldsymbol{K}_j^\star = \boldsymbol{K}^\star(\mathcal{G}_j)$ and using the matrix inversion lemma, we can reduce the update of $\boldsymbol{K}^\star$ to

$$\boldsymbol{K}^\star(\mathcal{G}_j\cup\{s\}) = \boldsymbol{K}_{j-1}^\star - \boldsymbol{W}\frac{\boldsymbol{K}_{j-1}^\star\boldsymbol{v}_u\boldsymbol{v}_u^H\boldsymbol{K}_{j-1}^\star}{\lambda_{w,u}^{-1} + \boldsymbol{v}_u^H\boldsymbol{K}_{j-1}^\star\boldsymbol{v}_u}. \tag{35}$$

Algorithm 2 Greedy sampling set selection

$\mathcal{G}_0 = \{\}$ and $\boldsymbol{K}_0^\star = \boldsymbol{\Lambda}$

for $j = 1, \ldots, \ell$

$u = \operatorname*{argmax}_{s\in\mathcal{V}\setminus\mathcal{G}_{j-1}} \dfrac{\boldsymbol{v}_u^H\boldsymbol{K}_{j-1}^\star\boldsymbol{W}\boldsymbol{K}_{j-1}^\star\boldsymbol{v}_u}{\lambda_{w,u}^{-1} + \boldsymbol{v}_u^H\boldsymbol{K}_{j-1}^\star\boldsymbol{v}_u}$ ▷ $\mathcal{O}(n|\mathcal{K}|^2)$

$\boldsymbol{K}_j^\star = \boldsymbol{K}_{j-1}^\star - \boldsymbol{W}\dfrac{\boldsymbol{K}_{j-1}^\star\boldsymbol{v}_u\boldsymbol{v}_u^H\boldsymbol{K}_{j-1}^\star}{\lambda_{w,u}^{-1} + \boldsymbol{v}_u^H\boldsymbol{K}_{j-1}^\star\boldsymbol{v}_u}$ ▷ $\mathcal{O}(|\mathcal{K}|^2)$

$\mathcal{G}_j = \mathcal{G}_{j-1}\cup\{u\}$

end

From linearity, it is then straightforward to see that finding the minimum of the trace of (35) is equivalent to finding the maximum of

$$\operatorname{Tr}\left[\boldsymbol{W}\frac{\boldsymbol{K}_{j-1}^\star\boldsymbol{v}_u\boldsymbol{v}_u^H\boldsymbol{K}_{j-1}^\star}{\lambda_{w,u}^{-1} + \boldsymbol{v}_u^H\boldsymbol{K}_{j-1}^\star\boldsymbol{v}_u}\right] = \frac{\boldsymbol{v}_u^H\boldsymbol{K}_{j-1}^\star\boldsymbol{W}\boldsymbol{K}_{j-1}^\star\boldsymbol{v}_u}{\lambda_{w,u}^{-1} + \boldsymbol{v}_u^H\boldsymbol{K}_{j-1}^\star\boldsymbol{v}_u}. \tag{36}$$

The greedy sampling set selection procedure obtained by leveraging (35) and (36) is presented in Algorithm 2. This algorithm now requires only $\mathcal{O}(n\ell|\mathcal{K}|^2)$ operations.

A. Simulations

In this section, we start by evaluating the performance greedy sampling set selection (Algorithm 2). For comparison, we also display the results obtained by the *uniform* and *leverage score* randomized methods from [6] and a *deterministic*

heuristics based on sampling nodes with the highest leverage score ($\|\boldsymbol{v}_i\|^2$). In the following examples, we use undirected graphs generated using the *Erdös-Rényi* model, in which an edge is placed between two nodes with probability $p = 0.2$; the *preferential attachment* model [36], in which nodes are added one at a time and connected to a node already in the graph with probability proportional to its degree; and a *random undirected graph*, obtained by assigning a weight to all possible edges uniformly at random from $[0, 1]$.

The figure of merit in the following simulations is the *relative suboptimality* from (26). Since it depends on the optimal sampling set which needs to be determined by exhaustive search, we first focus on graphs with $n = 20$ nodes. The bandlimited graph signals are generated by taking $\boldsymbol{V}_{\mathcal{K}}$ in (3) to be the eigenvectors of the graph adjacency matrix relative to the five eigenvalues with largest magnitude ($|\mathcal{K}| = 5$). The random vectors $\bar{\boldsymbol{x}}$ in (3) and $\boldsymbol{w}$ in (4) are realizations of zero-mean Gaussian random variables with covariance matrices $\boldsymbol{\Lambda} = \boldsymbol{I}$ and $\boldsymbol{\Lambda}_w = \sigma_w^2 \boldsymbol{I}$, where σ_w^2 is varied to obtain different SNRs. The transform in (6) is taken to be the identity ($\boldsymbol{H} = \boldsymbol{I}$) and the sampling set size is chosen as $\ell = |\mathcal{K}| = 5$.

Figures 7 and 8 display histograms of the relative suboptimality for 1000 realizations of graphs and graph signals for $\sigma_w^2 = 10^2$ (SNR $= -20$ dB) and $\sigma_w^2 = 10^{-2}$ (SNR $=$ 20 dB), respectively. As predicted by Theorems 2 and 4, greedy sampling set selection performs better in low SNR environments, where the optimal sampling set was obtained in more than 95% of the realizations. Nevertheless, even in high SNRs, it found the optimal sampling set almost half of the time. In fact, note that Algorithm 2 typically performs much better than the bounds in Theorems 2 and 4 (see details in Fig. 8). For illustration, results for greedily optimizing $\log\det\left[\boldsymbol{K}^\star(\mathcal{S})\right]$, a supermodular function, are shown in Figure 9. Although the MSE is α-supermodular with $\alpha < 1$, the relative suboptimality obtained by using both cost functions is similar.

It is worth noting that, although the deterministic leverage score ranking technique often yields good results, there are advantages to greedy sampling set selection, specially for higher SNR. The randomized sampling schemes, on the other hand, do not perform as well for single problem instances. To be fair, these methods are more appropriate when several sampling sets of the same graph signal are considered. Indeed, the performance measures in [6] hold in expectation over sampling realizations.

Evaluating the relative suboptimality for larger graphs is untractable. However, since these sampling set selection techniques build the sampling set sequentially, we can assess their performance in terms of the sampling set size required to obtain a given MSE reduction. Figure 10 displays the distribution of the sampling set size required to achieve a 90% reduction in the MSE with respect to the empty set. The plots are obtained from 1000 graphs and signals realizations with $n = 100$ nodes, $\boldsymbol{V}_{\mathcal{K}}$ in (3) composed the eigenvectors relative to the seven eigenvalues with largest magnitude ($|\mathcal{K}| = 7$), and $\sigma_w^2 = 10^{-2}$. Although Theorem 4 estimates that Algorithm 2 requires sets considerably larger to recover the same near-

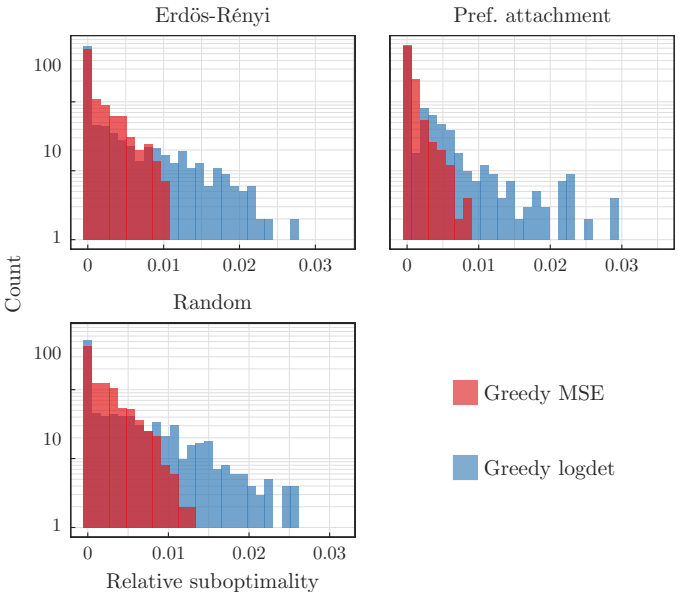

Figure 9. Relative suboptimality of MSE and log det (SNR $= 20$ dB)

optimal guarantees as supermodular functions, greedy sampling obtained a sampling set of size $|\mathcal{K}|$ in more than 50% of the realizations. Moreover, as noted in [6], we can now see that leverage score sampling has similar performance to uniform sampling for Erdös-Rényi graphs, but gives better results for the preferential attachment model.

B. Application: Subsampled Kernel PCA

Kernel PCA is a nonlinear version of PCA [24] that also identifies data subspaces by truncating the eigenvalue decomposition (EVD) of a Gram matrix $\boldsymbol{\Phi}$. However, whereas PCA uses the empirical covariance matrix, kPCA constructs $\boldsymbol{\Phi}$ by evaluating inner products between data points in a higher dimensional space $\mathbb{F}$ known as the *feature space*. Since the map $\varphi : \mathbb{R}^m \to \mathbb{F}$ can be nonlinear and $\mathbb{F}$ typically has infinite dimensionality, kPCA results in richer subspaces than PCA [24], [25], [37].

Naturally, the dimensionality of $\mathbb{F}$ poses a challenge for the construction of the Gram matrix. This problem is addressed using the so called *kernel trick* [24], [25], [37]. A kernel is a function κ that allows the inner product in $\mathbb{F}$ to be evaluated directly from vectors in $\mathbb{R}^m$, i.e., $\kappa(\boldsymbol{r}, \boldsymbol{s}) = \langle \varphi(\boldsymbol{r}), \varphi(\boldsymbol{s}) \rangle_{\mathbb{F}}$. We can use κ to construct $\boldsymbol{\Phi}$ from a training set $\{\boldsymbol{u}_i\}_{i=1,\ldots,n}$, $\boldsymbol{u}_i \in \mathbb{R}^m$ as in

$$\boldsymbol{\Phi} = \left[\kappa(\boldsymbol{u}_i, \boldsymbol{u}_j)\right]_{i,j=1,\ldots,n}. \tag{37}$$

Kernel PCA identifies the data subspace as the span of the first k eigenvectors of $\boldsymbol{\Phi}$, i.e., as $\operatorname{colspan}\{\boldsymbol{V}_{\mathcal{K}}\}$, where $\boldsymbol{\Phi} = \boldsymbol{V}\boldsymbol{\Lambda}\boldsymbol{V}^H$ is the EVD of $\boldsymbol{\Phi}$ with eigenvalues in decreasing order and $\mathcal{K} = 1, \ldots, k$. Using the representer's theorem, any data point $\boldsymbol{y}$ can be projected onto this subspace by

$$\bar{\boldsymbol{y}} = \boldsymbol{V}_{\mathcal{K}}^H \tilde{\boldsymbol{y}}, \quad \tilde{\boldsymbol{y}} = \left[\kappa(\boldsymbol{u}_i, \boldsymbol{y})\right]_{i=1,\ldots,n}. \tag{38}$$

The projection in (38) requires $\Theta(kn)$ operations and n KEs, making this method impractical for large data sets even

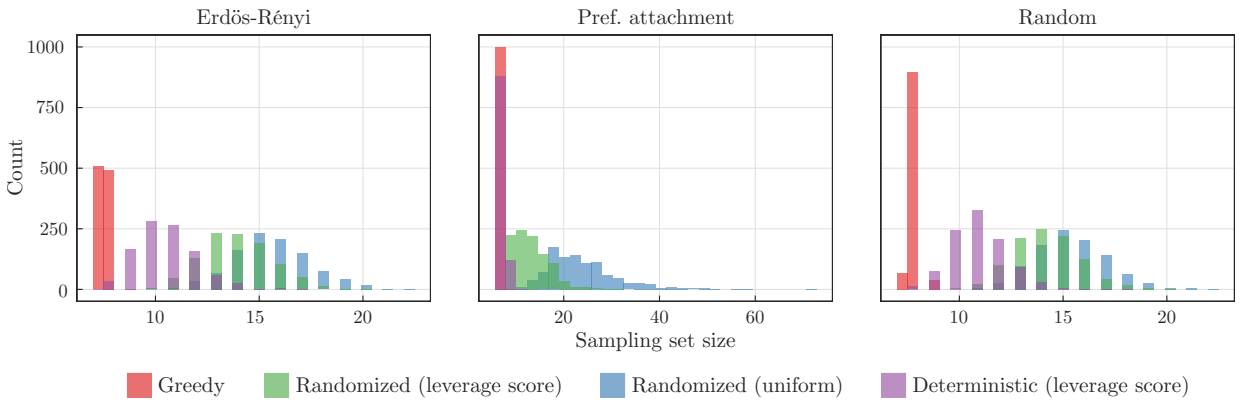

Figure 10. Sampling set size for 90% reduction of MSE

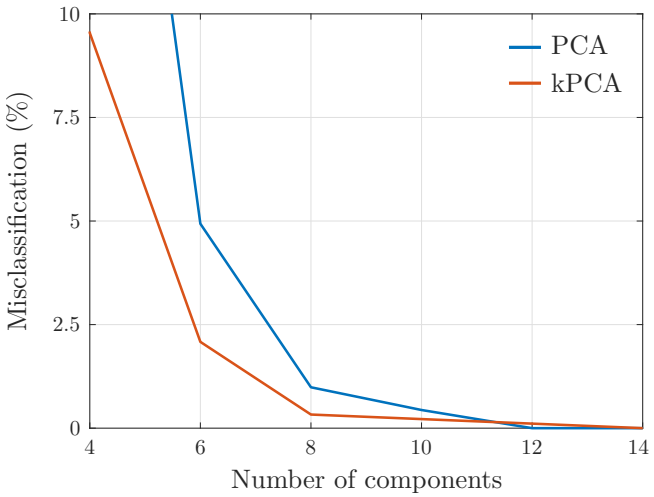

Figure 11. Classification performance of PCA and kPCA

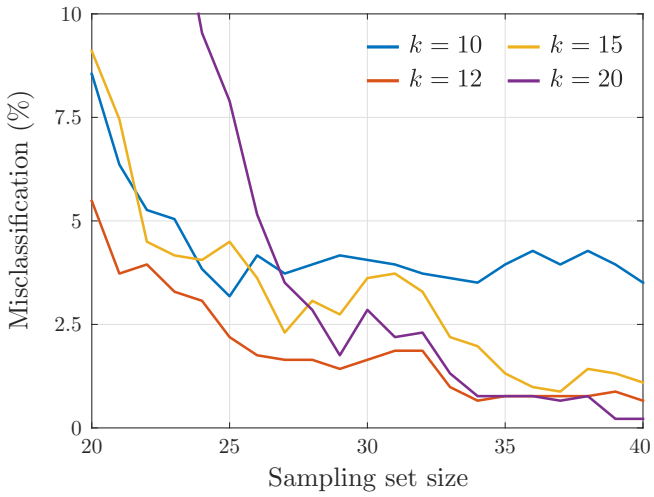

Figure 12. Classification performance of subsampled kPCA

if the dimension k of the subspace of interest is small. Indeed, although the training phase in (37) is usually performed offline, (38) needs to be evaluated during the operation phase for every new data point. In [38], this issue was addressed by using a Gaussian generative model for $\boldsymbol{\Phi}$ and showing that its maximum likelihood estimate depends only on a subset of the $\boldsymbol{u}_i$. Another approach is to impose sparsity on $\boldsymbol{V}$ *a priori* so that it depends only on a reduced number of training points [25]. Alternatively, one can find a representative subset of the training data and apply kPCA to that subset [39]. The issue with the latter method is that finding a good data subset is known to be a hard problem [22], [23]. In fact, it is related to the problem of sampling set selection in GSP.

Indeed, since we used the same notation as in Section II, formulating kPCA in the context of GSP is straightforward. Let the graph $\mathbb{G}$ have adjacency matrix $\boldsymbol{A} = \boldsymbol{\Phi}$, which is symmetric and normal, so that (38) has the form of a (partial) graph Fourier transform (2). In other words, (38) can be interpreted as enforcing graph signals of the form $\tilde{\boldsymbol{y}}$ to be bandlimited on $\boldsymbol{\Phi}$. Therefore, we can apply the sampling and interpolation theory from GSP to put forward a *subsampled kPCA*.

Based on the guarantees given in Section IV, we can use greedy search to obtain a sampling set $\mathcal{S}$ and use the interpolation techniques from Section II-A to recover $\tilde{\boldsymbol{y}}$ from its samples as in

$$\tilde{\boldsymbol{y}} = \boldsymbol{L}^{\star} \tilde{\boldsymbol{y}}_{\mathcal{S}}. \tag{39}$$

Then, (39) and (38) yield

$$\bar{\boldsymbol{y}} = \underbrace{\boldsymbol{V}_{\mathcal{K}}^{H} \boldsymbol{L}^{\star}}_{\boldsymbol{P}} \tilde{\boldsymbol{y}}_{\mathcal{S}}. \tag{40}$$

Notice that $\boldsymbol{P}$ is now $k \times |\mathcal{S}|$, so that the projection in (40) only takes $\Theta(k|\mathcal{S}|)$ operations and $|\mathcal{S}|$ KEs, leading to a considerable complexity reduction ($|\mathcal{S}|/n$) over the direct projection in (38). Moreover, kPCA is typically used for dimensionality reduction prior to regression or classification, so that we are interested in a linear transformation of $\bar{\boldsymbol{y}}$. Subsampled kPCA can account for this case by properly choosing $\boldsymbol{H}$ in (6). It is worth noting that contrary to [39], the full dataset is used during the training stage to obtain $\boldsymbol{V}_{\mathcal{K}}$. However, once $\boldsymbol{P}$ is determined, only the subset $\mathcal{S}$ is required.

In the sequel, we illustrate this method in a face recognition application using the *faces94* data set [40]. It contains 20 pictures of 152 individuals which were converted to black and white and normalized so that the value of each pixel is in $[-1, 1]$. A training set is obtained by randomly choosing 14

images for each individual (70% of the data set) and the remaining pictures are used for testing. In this application, we use a polynomial kernel of degree $d = 2$ [37] and a one-against-one multiclass support vector machine (SVM) classifier, in which an SVM is trained for each pair of class and the classification is obtained by majority voting (see [41] for details on this scheme). Finally, note that since images in both training and testing sets come from the same data set there is no observation noise $\boldsymbol{w}$. Still, σ_w^2 can be used to regularize the matrix inversions in (10) and (11) [26].

Figure 11 shows the misclassification percentage on the test set as a function of the number of components (k) for both PCA and kPCA. Note that kPCA can achieve the same performance as PCA with less components. The results of using the subsampled kPCA are shown in Figure 12 and clearly illustrate the trade-off between complexity and performance: as sampling set size increases, the classification errors decrease. However, since misclassification is a nonlinear function of the MSE, it may be advantageous to use more components instead of increasing the sampling set. For instance, kPCA requires $k = 7$ components to achieve a misclassification of 1%, so that evaluating the direct projection in (38) takes 2128 KEs and 27664 operations. Subsampled kPCA, on the other hand, can achieve the same performance with $k = 12$ components and $|\mathcal{S}| = 33$, i.e., 33 KEs and 759 operations, a complexity reduction of more than 97%. Nevertheless, using 7 components, subsampled kPCA would require $\mathcal{S}$ to almost be the full training set.

Remark 3. Although this section discussed kPCA, the same argument applies to the classical PCA. It is therefore straightforward to derive an analog *subsampled PCA* technique using (37)–(40).

VI. Conclusion

This work provided a solution to graph signal sampling problems by addressing the issue of sampling set selection in two ways. First, it derived universal bounds on the interpolation MSE which allow the quality of any sampling set or sampling heuristic to be evaluated by gauging how close their reconstruction performance is to the lower bound. Second, it provided near-optimality results for greedy MSE minimization, demonstrating that greedy sampling set selection is an effective sampling scheme, justifying its empirical success in the literature. To derive these results, we started by obtaining the optimal linear interpolation operator for noisy graph signal samples (Proposition 1). Based on this result, we derived universal bounds on the interpolation MSE (Theorem 1). Then, we showed that although the sampling set selection problem is combinatorial, its solution can be approximated using greedy search. By introducing the concept of approximate supermodularity, near-optimal guarantees were given for the reconstruction MSE (Theorems 2, 3, and 4), thus justifying the success of greedy sampling schemes. Simulations illustrated these results and applied them to reduce the complexity of kernel PCA projections.

Appendix A
Proof of Theorem 3

Once again, part (i) is a corollary of Lemma 1. To prove part (ii), we proceed as for Theorem 4. However, we now let $\boldsymbol{Z}(\mathcal{A}) = \boldsymbol{\Lambda}^{-1} + \sum_{i\in\mathcal{A}} \lambda_{w,i}^{-1} \boldsymbol{v}_i \boldsymbol{v}_i^H$ so that again the increment in (24) reads

$$f(\mathcal{A}\cup\{u\}) - f(\mathcal{A}) = \operatorname{Tr}\left[\boldsymbol{W}\left(\boldsymbol{Z}(\mathcal{A}) + \lambda_{w,u}^{-1}\boldsymbol{v}_u\boldsymbol{v}_u^H\right)^{-1} - \boldsymbol{W}\boldsymbol{Z}(\mathcal{A})^{-1}\right],$$

which using the matrix inversion simplifies to

$$f(\mathcal{A}\cup\{u\}) - f(\mathcal{A}) = -\operatorname{Tr}\left[\boldsymbol{W}\frac{\boldsymbol{Z}(\mathcal{A})^{-1}\boldsymbol{v}_u\boldsymbol{v}_u^H\boldsymbol{Z}(\mathcal{A})^{-1}}{\lambda_{w,u}^{-1} + \boldsymbol{v}_u^H\boldsymbol{Z}(\mathcal{A})^{-1}\boldsymbol{v}_u}\right] = -\frac{\boldsymbol{v}_u^H\boldsymbol{Z}(\mathcal{A})^{-1}\boldsymbol{W}\boldsymbol{Z}(\mathcal{A})^{-1}\boldsymbol{v}_u}{\lambda_{w,u}^{-1} + \boldsymbol{v}_u^H\boldsymbol{Z}(\mathcal{A})^{-1}\boldsymbol{v}_u}.$$

Using this expression, the expression for α in (24) becomes

$$\alpha = \min_{\substack{\mathcal{A}\subseteq\mathcal{B}\subseteq\mathcal{V}\\ v\notin\mathcal{B}}} \frac{\lambda_{w,u}^{-1} + \boldsymbol{v}_u^H\boldsymbol{Z}(\mathcal{B})^{-1}\boldsymbol{v}_u}{\lambda_{w,u}^{-1} + \boldsymbol{v}_u^H\boldsymbol{Z}(\mathcal{A})^{-1}\boldsymbol{v}_u} \frac{\boldsymbol{v}_u^H\boldsymbol{Z}(\mathcal{A})^{-1}\boldsymbol{W}\boldsymbol{Z}(\mathcal{A})^{-1}\boldsymbol{v}_u}{\boldsymbol{v}_u^H\boldsymbol{Z}(\mathcal{B})^{-1}\boldsymbol{W}\boldsymbol{Z}(\mathcal{B})^{-1}\boldsymbol{v}_u}. \tag{41}$$

Note that, although similar, (41) is not the same as (32).

We now bound (41), by noticing that for any set $\mathcal{X}\subseteq\mathcal{V}$

$$\mu_{\min} \le \lambda_{\min}\left[\boldsymbol{\Lambda}^{-1}\right] \le \lambda_{\min}\left[\boldsymbol{Z}(\mathcal{X})\right] \le \lambda_{\max}\left[\boldsymbol{Z}(\mathcal{X})\right] \le \lambda_{\max}\left[\boldsymbol{\Lambda}^{-1} + \boldsymbol{V}_{\mathcal{K}}^H\boldsymbol{\Lambda}_w^{-1}\boldsymbol{V}_{\mathcal{K}}\right] \le \mu_{\max}.$$

Thus, using the Rayleigh quotient inequalities, leads to

$$\alpha \ge \frac{\lambda_{w,u}^{-1} + \|\boldsymbol{v}_u\|_2^2\,\lambda_{\max}\left[\boldsymbol{Z}(\mathcal{B})\right]^{-1}}{\lambda_{w,u}^{-1} + \|\boldsymbol{v}_u\|_2^2\,\lambda_{\min}\left[\boldsymbol{Z}(\mathcal{A})\right]^{-1}} \frac{\lambda_{\min}\left[\boldsymbol{Z}(\mathcal{A})^{-1}\boldsymbol{W}\boldsymbol{Z}(\mathcal{A})^{-1}\right]}{\lambda_{\max}\left[\boldsymbol{Z}(\mathcal{B})^{-1}\boldsymbol{W}\boldsymbol{Z}(\mathcal{B})^{-1}\right]},$$

which can be simplified using the bounds on the eigenvalues of the product of PSD matrices from [35] to yield

$$\alpha \ge \frac{\lambda_{w,u}^{-1} + \|\boldsymbol{v}_u\|_2^2\,\mu_{\max}^{-1}}{\lambda_{w,u}^{-1} + \|\boldsymbol{v}_u\|_2^2\,\mu_{\min}^{-1}} \frac{\mu_{\max}^{-2}}{\kappa_2(\boldsymbol{W})\,\mu_{\min}^{-2}} \triangleq \alpha', \tag{42}$$

where $\kappa_2(\boldsymbol{W}) = \lambda_{\max}(\boldsymbol{W})/\lambda_{\min}(\boldsymbol{W})$ is the 2-norm condition number of $\boldsymbol{W}$. To obtain the expression in (**??**), notice that (42) is decreasing with respect to $\|\boldsymbol{v}_u\|_2^2$ and $\lambda_{w,u}^{-1}$. Indeed,

$$\frac{\partial\alpha'}{\partial\|\boldsymbol{v}_u\|_2^2} = \frac{\mu_{\max}^{-2}}{\kappa_2(\boldsymbol{W})\,\mu_{\min}^{-2}} \frac{\lambda_{w,u}^{-1}\left(\mu_{\max}^{-1} - \mu_{\min}^{-1}\right)}{\left(\lambda_{w,u}^{-1} + \|\boldsymbol{v}_u\|_2^2\,\mu_{\min}^{-1}\right)^2} \le 0$$

$$\frac{\partial\alpha'}{\partial\lambda_{w,u}} = \frac{\mu_{\max}^{-2}}{\kappa_2(\boldsymbol{W})\,\mu_{\min}^{-2}} \frac{\lambda_{w,u}^{-2}\,\|\boldsymbol{v}_u\|_2^2\left(\mu_{\max}^{-1} - \mu_{\min}^{-1}\right)}{\left(\lambda_{w,u}^{-1} + \|\boldsymbol{v}_u\|_2^2\,\mu_{\min}^{-1}\right)^2} \le 0$$

which are both non-positive because $0 < \mu_{\min} \le \mu_{\max}$ and $\kappa_2(\boldsymbol{W}) \ge 1$ [27]. Once again, we use the fact that $\|\boldsymbol{v}_u\|_2^2 \le 1$ to get the expression in (**??**).

References

[1] D. Shuman, S. Narang, P. Frossard, A. Ortega, and P. Vandergheynst, "The emerging field of signal processing on graphs: Extending high-dimensional data analysis to networks and other irregular domains," *IEEE Signal Process. Mag.*, vol. 30[3], pp. 83–98, 2013.

[2] A. Sandryhaila and J. Moura, "Discrete signal processing on graphs," *IEEE Trans. Signal Process.*, vol. 61[7], pp. 1644–1656, 2013.

[3] S. Narang and A. Ortega, "Perfect reconstruction two-channel wavelet filter banks for graph structured data," *IEEE Trans. Signal Process.*, vol. 60[6], pp. 2786–2799, 2012.
[4] N. Tremblay, G. Puy, R. Gribonval, and P. Vandergheynst, "Compressive spectral clustering," in *Int. Conf. on Mach. Learning*, 2016, pp. 1002—-1011.
[5] X. Zhu and M. Rabbat, "Approximating signals supported on graphs," in *Int. Conf. on Acoust., Speech and Signal Process.*, 2012, pp. 3921–3924.
[6] S. Chen, R. Varma, A. Singh, and J. Kovačević, "Signal recovery on graphs: Fundamental limits of sampling strategies," *IEEE Trans. Signal Process.*, vol. 2[4], pp. 539–554, 2016.
[7] A. Marques, S. Segarra, G. Leus, and A. Ribeiro, "Sampling of graph signals with successive local aggregations," *IEEE Trans. Signal Process.*, vol. 64[7], pp. 1832–1843, 2016.
[8] H. Shomorony and A. Avestimehr, "Sampling large data on graphs," in *Global Conf. on Signal and Inform. Process.*, 2014, pp. 933–936.
[9] M. Tsitsvero, S. Barbarossa, and P. Di Lorenzo, "Signals on graphs: Uncertainty principle and sampling," *IEEE Trans. Signal Process.*, vol. 64[18], pp. 4845–4860, 2016.
[10] A. Anis, A. Gadde, and A. Ortega, "Efficient sampling set selection for bandlimited graph signals using graph spectral proxies," *IEEE Trans. Signal Process.*, vol. 64[14], pp. 3775–3789, 2016.
[11] S. Chen, R. Varma, A. Sandryhaila, and J. Kovačević, "Discrete signal processing on graphs: Sampling theory," *IEEE Trans. Signal Process.*, vol. 63[24], pp. 6510–6523, 2015.
[12] S. Chepuri and G. Leus, "Subsampling for graph power spectrum estimation," in *Sensor Array and Multichannel Signal Process. Workshop*, 2016.
[13] A. Krause, A. Singh, and C. Guestrin, "Near-optimal sensor placements in Gaussian processes: Theory, efficient algorithms and empirical studies," *J. Mach. Learning Research*, vol. 9, pp. 235–284, 2008.
[14] A. Das and D. Kempe, "Submodular meets spectral: Greedy algorithms for subset selection, sparse approximation and dictionary selection," in *Int. Conf. on Mach. Learning*, 2011.
[15] G. Sagnol, "Approximation of a maximum-submodular-coverage problem involving spectral functions, with application to experimental designs," *Discrete Appl. Math.*, vol. 161[1-2], pp. 258–276, 2013.
[16] J. Ranieri, A. Chebira, and M. Vetterli, "Near-optimal sensor placement for linear inverse problems," *IEEE Trans. Signal Process.*, vol. 62[5], pp. 1135–1146, 2014.
[17] F. Gama, A. Marques, G. Mateos, and A. Ribeiro, "Rethinking sketching as sampling: Linear transforms of graph signals," in *Asilomar Conf. on Signals, Syst. and Comput.*, 2016.
[18] D. Thanou, D. Shuman, and P. Frossard, "Learning parametric dictionaries for signals on graphs," *IEEE Trans. Signal Process.*, vol. 62[15], pp. 3849–3862, 2014.
[19] B. Mirzasoleiman, A. Badanidiyuru, A. Karbasi, J. Vondrák, and A. Krause, "Lazier than lazy greedy," in *AAAI Conf. on Artificial Intell.*, 2015, pp. 1812–1818.
[20] F. Bach, "Learning with submodular functions: A convex optimization perspective," *Foundations and Trends in Machine Learning*, vol. 6[2-3], pp. 145–373, 2013.
[21] L. Chamon and A. Ribeiro, "Near-optimality of greedy set selection in the sampling of graph signals," in *Global Conf. on Signal and Inform. Process.*, 2016.
[22] D. Woodruff, "Sketching as a tool for numerical linear algebra," *Foundations and Trends in Theoretical Computer Science*, vol. 10[1-2], pp. 1–157, 2014.
[23] D. Feldman, M. Schmidt, and C. Sohler, "Turning big data into tiny data: Constant-size coresets for K-means, PCA and projective clustering," in *ACM-SIAM Symp. on Discrete Algorithms*, 2013, pp. 1434–1453.
[24] B. Schölkopf, , A. Smola, and K.-R. Müller, "Nonlinear component analysis as a kernel eigenvalue problem," *Neural Comput.*, vol. 10[5], pp. 1299–1319, 1998.
[25] J. Arenas-Garcia, K. Petersen, G. Camps-Valls, and L. Hansen, "Kernel multivariate analysis framework for supervised subspace learning: A tutorial on linear and kernel multivariate methods," *IEEE Signal Process. Mag.*, vol. 30[4], pp. 16–29, 2013.
[26] T. Kailath, A. Sayed, and B. Hassibi, *Linear estimation*. Prentice-Hall, 2000.
[27] R. Horn and C. Johnson, *Matrix analysis*. Cambridge University Press, 2013.
[28] B. Girault, "Stationary graph signals using an isometric graph translation," in *European Signal Process. Conf.*, 2015, pp. 1516–1520.
[29] A. Marques, S. Segarra, G. Leus, and A. Ribeiro, "Stationary graph processes and spectral estimation," 2016, arXiv:1603.04667v1.
[30] N. Perraudin and P. Vandergheynst, "Stationary signal processing on graphs," 2016, arXiv:1601.02522v3.
[31] L. Chamon and A. Ribeiro, "Universal bounds for the sampling of graph signals," in *Int. Conf. on Acoust., Speech and Signal Process.*, 2017.
[32] R. Bhatia, *Matrix analysis*. Springer, 1997.
[33] G. Nemhauser, L. Wolsey, and M. Fisher, "An analysis of approximations for maximizing submodular set functions—I," *Mathematical Programming*, vol. 14[1], pp. 265–294, 1978.
[34] S. Joshi and S. Boyd, "Sensor selection via convex optimization," *IEEE Trans. Signal Process.*, vol. 57[2], pp. 451–462, 2009.
[35] B. Wang and F. Zhang, "Some inequalities for the eigenvalues of the product of positive semidefinite Hermitian matrices," *Linear Algebra and its Applications*, vol. 160, pp. 113–118, 1992.
[36] A.-L. Barabási and R. Albert, "Emergence of scaling in random networks," *Science*, vol. 286[5439], pp. 509–512, 1999.
[37] C. Bishop, *Pattern recognition and machine learning*. Springer, 2007.
[38] M. Tipping, "Sparse kernel principal component analysis," in *NIPS*, 2000.
[39] Y. Washizawa, "Subset kernel principal component analysis," in *Int. Workshop on Mach. Learning for Signal Process.*, 2009.
[40] L. Spacek, "Collection of facial images," http://cswww.essex.ac.uk/mv/allfaces/index.html.
[41] C.-W. Hsu and C.-J. Lin, "A comparison of methods for multiclass support vector machines," *IEEE Trans. Neural Netw.*, vol. 13[2], pp. 415–425, 2002.